Rethinking Sketching as Sampling: A Graph Signal Processing Approach

Fernando Gama, Antonio G. Marques, Gonzalo Mateos, and Alejandro Ribeiro

arXiv:1611.00119v1 [cs.IT] 1 Nov 2016

***Abstract*—Sampling of bandlimited graph signals has well-documented merits for dimensionality reduction, affordable storage, and online processing of streaming network data. Most existing sampling methods are designed to minimize the error incurred when reconstructing the original signal from its samples. Oftentimes these parsimonious signals serve as inputs to computationally-intensive linear operator (e.g., graph filters and transforms). Hence, interest shifts from reconstructing the signal itself towards instead approximating the output of the prescribed linear operator efficiently. In this context, we propose a novel sampling scheme that leverages the bandlimitedness of the input as well as the transformation whose output we wish to approximate. We formulate problems to jointly optimize sample selection and a sketch of the target linear transformation, so when the latter is affordably applied to the sampled input signal the result is close to the desired output. These designs are carried out off line, and several heuristic (sub)optimal solvers are proposed to accommodate high-dimensional problems, especially when computational resources are at a premium. Similar sketching as sampling ideas are also shown effective in the context of linear inverse problems. The developed sampling plus reduced-complexity processing pipeline is particularly useful for streaming data, where the linear transform has to be applied fast and repeatedly to successive inputs or response signals. Numerical tests show the effectiveness of the proposed algorithms in classifying handwritten digits from as few as 20 out of 784 pixels in the input images, as well as in accurately estimating the frequency components of bandlimited graph signals sampled at few nodes.**

***Index Terms*—Sketching, sampling, graph signal processing, streaming, linear transforms, linear inverse problems**

I. Introduction

Propelled by the desire of analyzing and processing data supported on irregular domains, there has been a growing interest in broadening the scope of traditional signal processing techniques to signals defined on graphs [3], [4]. Noteworthy representatives include sampling of graph signals, linear graph filtering and the graph Fourier transform (GFT) [5], [6], all of them instances of linear problems. This is not surprising since linear models are ubiquitous in science and engineering, due in part to their generality, conceptual simplicity, and mathematical tractability. Along with heterogeneity and lack of regularity, data are increasingly high dimensional and this curse of dimensionality not only raises statistical challenges, but also major computational hurdles even for linear models [7]. In particular, these limiting factors can hinder processing of streaming data, where say a massive linear operator has to be repeatedly and efficiently applied to a sequence of input (graph) signals [8]. These Big Data challenges motivated a recent body of work collectively addressing so-termed *sketching* problems [9], [10], which seek computationally-efficient solutions to a subset of (typically inverse) linear problems. The basic idea is to *draw a sketch* of the linear model such that the resulting linear transform is lower dimensional, while still offering quantifiable approximation error guarantees. To achieve this, a fat random projection matrix is designed to pre-multiply and reduce the dimensionality of the linear operator matrix, in such way that the resulting matrix sketch still captures the quintessential structure of the model. The input vector has to be adapted to the sketched operator as well, and to that end the same random projections are applied to the signal in a way often agnostic to the input statistics.

Although random projection methods offer an elegant dimensionality-reduction alternative for several Big Data problems, they face some shortcomings: i) sketching each new input signal entails a nontrivial computational cost, which can be a bottleneck in streaming applications; ii) the design of the random projection matrix does not take into account any a priori information on the input; and iii) the guarantees offered are probabilistic. Alternatively one can think of reducing complexity by simply retaining a few samples of each input. Different from random sketching, under stationarity the sampling can remain fixed and it incurs negligible online complexity. Sampling schemes typically build on a parsimonious model for the signals of interest, which in the case of graph signals is either smoothness or bandlimitedness – i.e., a sparse representation in the graph Fourier domain [11]–[13]. However, most existing sampling methods are designed with the objective of reconstructing the original graph signal, and do not account for subsequent processing the signal may undergo; see [14]–[16] for exceptions, the latter two dealing with graph signals.

In this sketching context and towards reducing the online computational cost of obtaining the solution to a linear problem, we propose a novel sampling scheme for signals that belong to a *known* subspace. Since such signals can be conveniently represented as bandlimited graph signals (see Remark 1), we leverage graph signal processing results to build our online schemes. Different from most existing sampling approaches, our design explicitly accounts for the transformation whose output we wish to approximate. The approach is

Work in this paper is supported by USA NSF CCF-1217963 and Spanish MINECO grant No TEC2013- 41604-R. F. Gama and A. Ribeiro are with the Dept. of Electrical and Systems Eng., Univ. of Pennsylvania., A. G. Marques is with the Dept. of Signal Theory and Comms., King Juan Carlos Univ., G. Mateos is with the Dept. of Electrical and Computer Eng., Univ. of Rochester. Emails: fgama@seas.upenn.edu, antonio.garcia.marques@urjc.es, gmateosb@ece.rochester.edu, and aribeiro@seas.upenn.edu. Part of the results in this paper were submitted to the *2016 Asilomar Conf. of Signals, Systems and Computers* [1] and the *2016 IEEE GobalSIP Conf.* [2].

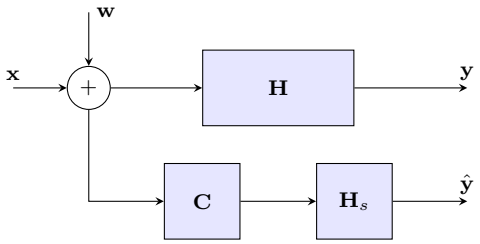

Figure 1. Direct sketching problem. Knowing the linear transform matrix $\mathbf{H}$ and having access to a stream of noisy inputs $(\mathbf{x}+\mathbf{w})$, we want to estimate the output $\mathbf{y}$ of this system in a fast way, reducing the online computational cost by operating with the matrix sketch $\mathbf{H}_s$ only on p samples selected by the selection matrix $\mathbf{C}$, $\hat{\mathbf{y}} = \mathbf{H}_s\mathbf{C}(\mathbf{x}+\mathbf{w})$.

then to shift the computational burden to the off-line phase where both the sampling pattern and the sketch of the linear transformation are designed. After doing that once, the online phase merely consists in selecting the signal values dictated by the sampling pattern and process these samples using the sketch of the linear transformation.

In Section II we introduce the mathematical formulation of the direct and inverse linear problems as well as the assumptions on the input (graph) signals. Then, we proceed to present the solutions for the direct and inverse problems in Sections III and IV, respectively. In both cases, we obtain first a closed form expression for the optimal reduced linear transform as a function of the selected samples of the signal. Then we use that expression to obtain an equivalent optimization problem on the selection of samples, that turns out to be a Semidefinite Program (SDP) modulo binary constraints that arise naturally from the sample (node) selection problem. Section V discusses a number of heuristics to obtain tractable solutions to the binary optimization. In Section VI we apply this framework to the problem of estimating the graph frequency components of a graph signal in a fast and efficient fashion as well as the problem of classifying handwritten digits. Finally, conclusions are drawn in Section VII.

Notation: Generically, the entries of a matrix $\mathbf{X}$ and a (column) vector $\mathbf{x}$ will be denoted as X_{ij} and x_i. The notation ^{T} and ^{H} stands for transpose and transpose conjugate, respectively, and the superscript † denotes pseudoinverse; $\mathbf{0}$ is the all-zero vector and $\mathbf{1}$ is the all-one vector; and the ℓ_0 pseudo norm $\|\mathbf{X}\|_0$ equals the number of nonzero entries in $\mathbf{X}$. For a vector $\mathbf{x}$, $\text{diag}(\mathbf{x})$ is a diagonal matrix with the (i,i)th entry equal to x_i; when applied to a matrix, $\text{diag}(\mathbf{X})$ is a vector with the diagonal elements of $\mathbf{X}$.

II. Problem formulation

We consider here linear sketching problems for a collection of signals $\{\mathbf{x}_t\}_{t\in\mathbb{N}}$ that are assumed to belong to a linear subspace or, equivalently, that are assumed to be bandlimited in a graph. The overall idea is to regard $\{\mathbf{x}_t\}_{t\in\mathbb{N}}$ as realizations of a given random process and leverage the bandlimited assumption to reduce the number of samples required to process the signals of interest. For a precise problem formulation let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ be a graph described by a set of n nodes $\mathcal{V}$, a set $\mathcal{E}$ of edges (i,j) and a weight function $\mathcal{W} : \mathcal{E} \to \mathbb{R}$ that assigns weights to the directed edges. A signal $\mathbf{x} \in \mathbb{R}^n$ can be defined on the nodes of such a network where each element of the vector $\mathbf{x} = [x_1, \ldots, x_n]^T$ represents a real value present at the node [3], [17]. In order to capture the structure of the network and its impact on the signal, a graph-shift operator $\mathbf{S} \in \mathbb{R}^{n\times n}$ is introduced [18]. Matrix $\mathbf{S}$ is such that $[\mathbf{S}]_{i,j} = 0$ whenever $i \neq j$, $(j,i) \notin \mathcal{E}$. The graph shift operator here is considered normal so that we can write

$$\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^H = \mathbf{V}\text{diag}\left([\lambda_1, \ldots, \lambda_n]^T\right)\mathbf{V}^H, \tag{1}$$

where $\mathbf{V} = [\mathbf{v}_1, \ldots, \mathbf{v}_n] \in \mathbb{C}^{n\times n}$ in (1) contains the eigenvectors of $\mathbf{S}$ and the diagonal matrix $\mathbf{\Lambda} = \text{diag}\left([\lambda_1, \ldots, \lambda_n]^T\right) \in \mathbb{C}^{n\times n}$ the corresponding eigenvalues. Since the matrix $\mathbf{V}$ is unitary it holds $\mathbf{V}^{-1} = \mathbf{V}^H$. Examples of normal graph shift operators are the adjacency matrices of some graphs and the Laplacian of any undirected graph, and their respective normalized counterparts [19]. The shift operator $\mathbf{S}$ is used to define the Graph Fourier Transform (GFT) $\tilde{\mathbf{x}} = \mathbf{V}^H\mathbf{x}$ and its inverse $\mathbf{x} = \mathbf{V}\tilde{\mathbf{x}}$; see, e.g., [6]. The inverse (i)GFT is clearly a proper inverse since we can always write

$$\mathbf{x} \;=\; \mathbf{V}\tilde{\mathbf{x}} \;=\; \mathbf{V}\left(\mathbf{V}^H\mathbf{x}\right). \tag{2}$$

The elements of the vector $\tilde{\mathbf{x}} = [\tilde{x}_1, \ldots, \tilde{x}_n]^T$ are said to be the frequency coefficients of the signal $\mathbf{x}$ and we think of the iGFT relationship $\mathbf{x} = \mathbf{V}\tilde{\mathbf{x}}$ as an alternative representation of $\mathbf{x}$ tailored to the structure of the shift operator $\mathbf{S}$.

Throughout this paper we model the signals $\{\mathbf{x}_t\}_{t\in\mathbb{N}}$ as realizations of a zero-mean random process $\mathbf{x}$ which is *bandlimited* in $\mathbf{S}$. Formally, we assume that there exists a constant $k \ll n$ such that $\tilde{x}_{k'} = 0$ for all $k' > k$. This means that we can define the vector $\tilde{\mathbf{x}}_k = [\tilde{x}_1, \ldots, \tilde{x}_k]^T$ containing the first k elements of $\tilde{\mathbf{x}}$ to write the GFT $\tilde{\mathbf{x}} = [\tilde{\mathbf{x}}_k^T, \mathbf{0}_{n-k}^T]^T$ as the stacking of $\tilde{\mathbf{x}}_k$ and the all-zero vector of size $n-k$. If this is the case, we can group the first k eigenvectors of $\mathbf{S}$ in the matrix $\mathbf{V}_k = [\mathbf{v}_1, \ldots, \mathbf{v}_k] \in \mathbb{C}^{n\times k}$ and rewrite (2) as

$$\mathbf{x} \;=\; \mathbf{V}_k\tilde{\mathbf{x}}_k \;=\; [\mathbf{V}_k, \mathbf{0}_{n-k}]\tilde{\mathbf{x}} \;=\; [\mathbf{V}_k, \mathbf{0}_{n-k}]\left(\mathbf{V}^H\mathbf{x}\right). \tag{3}$$

The bandlimitedness of $\mathbf{x}$ also shows up in its statistical description. To be rigorous, let the covariance matrix of $\mathbf{x}$ be denoted by $\mathbf{R}_x = \mathbb{E}[\mathbf{x}\mathbf{x}^T] \in \mathbb{R}^{n\times n}$. Then, we can use (3) to rewrite $\mathbf{R}_x = \mathbf{V}_k\mathbb{E}[\tilde{\mathbf{x}}_k\tilde{\mathbf{x}}_k^H]\mathbf{V}_k^H = \mathbf{V}_k\mathbf{T}\mathbf{V}_k^H$ where $\mathbf{T} \in \mathbb{C}^{k\times k}$ is the so-termed frequency template defined as

$$\mathbf{T} \;=\; \mathbb{E}\left[\tilde{\mathbf{x}}_k\tilde{\mathbf{x}}_k^H\right]. \tag{4}$$

If no information is available about the process $\mathbf{x}$ other than it being bandlimited, one can assume $\tilde{\mathbf{x}}_k$ to be white, so that the template is $\mathbf{T} = \mathbf{I}_k$, i.e., the identity matrix of size $k \times k$. Identity and diagonal templates arise naturally if we assume that $\mathbf{x}$ is stationary with respect to the shift operator $\mathbf{S}$; see [20]–[22]. Regardless of the color of the template, it follows that the symmetric positive semidefinite matrix $\mathbf{R}_x$ is *singular*, because the rank of $\mathbf{T}$ (hence, the rank of $\mathbf{R}_x$) is at most k.

Another key implication of $\mathbf{x}$ being bandlimited in $\mathbf{S}$ is that it can be perfectly recovered from k samples. These k samples can be a selection of values at k nodes [11], a selection of values following from the local aggregation of a diffusion

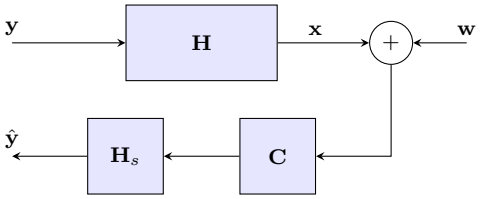

Figure 2. Inverse sketching problem. Given knowledge of the linear transform matrix $\mathbf{H}$ and having access to a stream of noisy observations of the output $(\mathbf{x}+\mathbf{w})$, we want to estimate the input $\mathbf{y}$ generating such outputs in a ways that minimizes the online computational cost. This is achieved by designing a matrix sketch $\mathbf{H}_s$ that operates on selected samples of the measured output $\mathbf{C}(\mathbf{x}+\mathbf{w})$.

process on $\mathbf{S}$ at a single node [12], or a combination of both strategies [23]. In either case, a parsimonious equivalent description of $\mathbf{x}$ becomes available and can be used to process $\mathbf{x}$ with a much smaller computational cost if $k \ll n$. In this paper we exploit this observation to reduce the computational cost of implementing a very fat linear transformation and to reduce the computational cost of implementing a very tall inverse linear transformation. We introduce these problems in the following section.

Remark 1 (**Bandlimitedness and subspace models**)**.** While (3) is derived under the assumption that $\mathbf{x}$ is a bandlimited graph signal, the expression holds true for *any* signal $\mathbf{x}$ belonging to a signal subspace $\mathcal{X}_k$ of dimension k. In fact, with $\mathbf{V}_k$ being an orthonormal basis of $\mathcal{X}_k$ and defining the matrix $\mathbf{P}_{\mathcal{X}_k} = \mathbf{V}_k\mathbf{V}_k^H$, which projects any input signal into $\mathcal{X}_k$, it follows readily that (3) can be written as $\mathbf{x} = \mathbf{P}_{\mathcal{X}_k}\mathbf{x}$, demonstrating that our bandlimited signals can be equivalently thought as signals lying on $\mathcal{X}_k$. Although the exposition in this section assumed that the graph shift operator $\mathbf{S}$ was known and used $\mathbf{S}$ to find $\mathbf{V}_k$ and $\mathcal{X}_k$, one can also follow the reverse path and use $\mathcal{X}_k$ to find first $\mathbf{V}_k$ and then $\mathbf{S}$. Schemes to identify $\mathbf{S}$ based on $\mathbf{V}_k$ have been reported in [24], [25].

A. Linear sketching problems

We start by formulating the optimal direct and inverse sketching problems and then discuss a suboptimal design.

1) Direct sketching problem: The setup of the *direct* linear sketching problem is illustrated in Fig. 1. We are given a known matrix $\mathbf{H} \in \mathbb{R}^{m\times n}$ and observe a possibly noisy realization of a signal $\mathbf{x}$ that is bandlimited in the graph shift $\mathbf{S}$ with spectral template $\mathbf{T}$. The goal is to compute an estimate of the output of the linear transformation $\mathbf{y} = \mathbf{H}\mathbf{x}$. As an interesting aside, the fact of $\mathbf{x}$ being bandlimited in $\mathbf{S}$ dictates that the optimal transformation can be written as a low-pass filter on the graph $\mathbf{S}$, followed by a lower dimensional linear transformation. Assuming that this two steps are combined into a single one, their computational cost is of order $O(mn)$. Our operating assumption, however, implies that it is possible to represent $\mathbf{x}$ with $k \ll n$ samples. We therefore proceed through the alternative route of first drawing $p \geq k$ samples from the signal $\mathbf{x}$ to yield the sampled signal $\mathbf{C}\mathbf{x}$, where $\mathbf{C}$ is a $p \times n$ sampling matrix that selects p out the n elements of $\mathbf{x}$. To process the reduced-dimensionality sampled input, the goal is to optimally design a sketch $\mathbf{H}_s \in \mathbb{R}^{m\times p}$ of the linear transform $\mathbf{H}$, so that $\hat{\mathbf{y}} = \mathbf{H}_s\mathbf{C}\mathbf{x}$ is an accurate estimate of the intended transformation $\mathbf{y} = \mathbf{H}\mathbf{x}$. Given $\mathbf{H}_s$ and $\mathbf{C}$, the cost of forming $\hat{\mathbf{y}}$ is of order $O(mp)$ because the sampled vector $\mathbf{C}\mathbf{x}$ contains p elements only. The computational cost is then reduced by a factor p/n, which can be substantial if we select $k \leq p \ll n$.

For a formal problem statement let us begin by defining the set $\mathcal{C}_{pn}$ of admissible $p \times n$ sampling matrices $\mathbf{C}$, for given p and n. Such matrices must have binary elements $C_{ij} \in \{0,1\}$ and be such that each row contains *exactly* one nonzero element while each column contains *at most* one nonzero element. Accordingly, the set of all matrices of this form is given by

$$\mathcal{C}_{pn} := \{\mathbf{C} \in \{0,1\}^{p\times n} : \ \mathbf{C}\mathbf{1} = \mathbf{1}, \ \mathbf{C}^T\mathbf{1} \preceq \mathbf{1}, \ \}. \quad (5)$$

For upcoming derivations it is important to observe that selection matrices $\mathbf{C} \in \mathcal{C}_{pn}$ have rank p and satisfy $\mathbf{C}\mathbf{C}^T = \mathbf{I}_p$. The selection matrix also satisfies $\mathbf{C}^T\mathbf{C} = \text{diag}(\mathbf{c})$, where $\mathbf{c}$ is a binary vector of size n, whith $c_i = 1$ if and only if the i-th column of $\mathbf{C}$ contains a one, or, equivalently, if the i-th element of $\mathbf{x}$ is sampled. There is a one-to-one correspondence between the sampling matrix $\mathbf{C}$ and the sampling vector $\mathbf{c}$, modulo the ordering of the rows in $\mathbf{C}$.

When the input signal is corrupted by additive noise $\mathbf{w}$ the sampled signal is $\mathbf{C}(\mathbf{x}+\mathbf{w})$; see Fig. 1. After processing the samples with the linear transformation $\mathbf{H}_s$, we obtain the estimate $\hat{\mathbf{y}} = \mathbf{H}_s\mathbf{C}(\mathbf{x}+\mathbf{w})$. Our goal here is to find a sampling matrix $\mathbf{C}$ and a sketch $\mathbf{H}_s \in \mathbb{R}^{m\times p}$ that minimizes the mean-squared error (MSE) with respect to the desired response $\mathbf{y} = \mathbf{H}\mathbf{x}$. We are therefore searching for the matrices $\mathbf{C}^*$ and $\mathbf{H}_s^*$ that solve the non-convex optimization problem

$$\{\mathbf{C}^*, \mathbf{H}_s^*\} := \underset{\mathbf{C}\in\mathcal{C}_{pn},\mathbf{H}_s}{\text{argmin}} \ \mathbb{E}\left[\left\|\mathbf{H}_s\mathbf{C}(\mathbf{x}+\mathbf{w}) - \mathbf{H}\mathbf{x}\right\|_2^2\right]. \quad (6)$$

We emphasize here that solving (6) is costly, indeed, intractable due to the binary constraints $\mathbf{C} \in \mathcal{C}_{pn}$. Methods for approximately solving (6) with manageable cost are presented in Section V, but these approaches still incur a cost much higher than just evaluating the linear transformation $\mathbf{y} = \mathbf{H}\mathbf{x}$. As pointed out in the introduction, the operating assumption in this work is that the optimization in (6) is carried out off line and only once, whereas during online operation we are presented with a stream of bandlimited signals $\{\mathbf{x}_t\}_{t\in\mathbb{N}}$. While the original goal is to process each of those signals with the transformation $\mathbf{H}$ [26], through the proposed sampling plus reduced-complexity processing pipeline, for every t we first sample $\mathbf{x}_t$ according to $\mathbf{C}^*$ and then apply the sketch $\mathbf{H}_s^*$ to the sampled signal $\mathbf{C}^*\mathbf{x}_t$, reducing the online operational cost by a factor of p/n. See Section III for a thorough treatment of the direct problem.

2) Inverse sketching problem: The basic setting for the related linear *inverse* sketching problem is illustrated in Fig. 2. We are again given a matrix $\mathbf{H} \in \mathbb{R}^{n\times m}$ and we have access to possibly noisy realizations of the output $\mathbf{x}$ of this system. These realizations are also bandlimited in the graph shift $\mathbf{S}$ and have a spectral template $\mathbf{T}$. In this case, we aim at estimating the input $\mathbf{y}$ that generated the observed output $\mathbf{x} = \mathbf{H}\mathbf{y}$. If $\mathbf{H}$

is full rank, then the ordinary least squares (LS) estimator of $\mathbf{y}$ is given by

$$\hat{\mathbf{y}}_{\rm LS} = \mathbf{A}_{\rm LS}\mathbf{x}, \quad \text{with } \mathbf{A}_{\rm LS} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T. \tag{7}$$

Computing matrix $\mathbf{A}_{\rm LS}$ demands $O(mn^2)$ calculations and its application to the output in order to estimate the input entails additional mn operations [27]. Under the assumption that $\mathbf{x}$ is bandlimited and can thus be represented with $k \ll n$ samples, an alternative procedure is to first design an appropriate sketch of $\mathbf{x}$ and matrix $\mathbf{A}_{\rm LS}$. This way, an estimate $\hat{\mathbf{y}}$ can be obtained by sampling the output and applying it to a sketch $\mathbf{H}_s \in \mathbb{R}^{m\times p}$ of the inverse mapping. More formally, given noisy observations of the output $\mathbf{x}+\mathbf{w}$, we want to form an estimator of the input $\hat{\mathbf{y}} = \mathbf{H}_s\mathbf{C}(\mathbf{x}+\mathbf{w})$ that minimizes the MSE between the predicted response $\mathbf{H}\hat{\mathbf{y}}$ and the actual one $\mathbf{x} = \mathbf{H}\mathbf{y}$. The optimun pair $\{\mathbf{C}^*, \mathbf{H}_s^*\}$ is thus given by

$$\{\mathbf{C}^*, \mathbf{H}_s^*\} := \underset{\mathbf{C}\in\mathcal{C}_{pn}, \mathbf{H}_s}{\operatorname{argmin}} \mathbb{E}\left[\left\|\mathbf{H}\mathbf{H}_s\mathbf{C}(\mathbf{x}+\mathbf{w}) - \mathbf{x}\right\|_2^2\right]. \tag{8}$$

Again, solving (8) is in general intractable, and even natural relaxations can be computationally expensive. However, recall that the streaming setting considered here calls for repetitive application of the inverse mapping $\mathbf{A}_{\rm LS}$, a task which incurs $O(mn)$ cost per datum. Thus, having available a sketch $\mathbf{H}_s$ that is applied only to a subset of p samples of $\mathbf{x}$ speeds up the online computation of the solution by a factor of $O(n/p)$. This problem is addressed in Section IV.

3) Operator sketching via row or column sampling: There may be applications where the sketch $\mathbf{H}_s$ must obey additional constraints. This can be the case in distributed setups where the values of $\mathbf{H}$ cannot be adapted or when the calculation (or storage) of the optimal $\mathbf{H}_s$ is impossible. In those setups, a design of a suboptimal sketch $\mathbf{H}_s$ based directly on the original values of $\mathbf{H}$ – e.g., via sampling of its rows or columns – may be the only alternative. To be more specific, consider first the case of the direct problem where $\mathbf{y} = \mathbf{H}\mathbf{x}$. Then, selecting $\mathbf{C}$ will not only determine the reduced input as $\mathbf{C}\mathbf{x}$ but also the sketch as $\mathbf{H}_s = \mathbf{H}\mathbf{C}^T$ – i.e., the columns of $\mathbf{H}$ corresponding to the selected samples –, so that the estimated output will be $\hat{\mathbf{y}} = \mathbf{H}\mathbf{C}^T\mathbf{C}\mathbf{x}$. Likewise in the inverse problem, sampling the output $\mathbf{x}$ amounts also to selecting rows of $\mathbf{H}$. Although these more constrained sampling-sketching designs will give rise to a lower performance, they can be nevertheless optimized. The details can be found in Sections III-B and IV-B.

B. Sketching and sampling

The problems explained in the previous section are related to both sampling and (traditional) matrix sketching [9], [10]. Once those problems have been formally introduced, the purpose of this section is to elaborate on the existing relation.

In the context of linear inverse problems, the objective of traditional sketching techniques is to reduce the dimension of the linear transformation $\mathbf{H}$ so that the number of computations required for obtaining an estimate of $\mathbf{y}$ is reduced [9], [10]. The sketching matrix $\mathbf{K} \in \mathbb{R}^{p\times n}$, with $p \ll n$, is designed such that the compressed quantities $\mathbf{K}\mathbf{H} \in \mathbb{R}^{p\times m}$ and $\mathbf{K}\mathbf{x} \in \mathbb{R}^p$ can be calculated fast. Given $\mathbf{K}$, an approximate solution to the (inverse) problem can be found as

$$\hat{\mathbf{y}}_{\rm S} = \min_{\mathbf{y}} \|(\mathbf{K}\mathbf{H})\mathbf{y} - (\mathbf{K}\mathbf{x})\|_2^2,$$

whose solution can be written in closed form as $\hat{\mathbf{y}}_{\rm S} = (\mathbf{K}\mathbf{H}\mathbf{H}^H\mathbf{K}^H)^{-1}\mathbf{K}\mathbf{x}$. The sketching matrix $\mathbf{K}$ is (usually) random and need not be sparse, although it must have a special structure that facilitates its efficient multiplication with $\mathbf{H}$ and $\mathbf{x}$. It is important to observe that within this approach there is no particular distinction between the off-line cost of computing $\mathbf{K}\mathbf{H}$ and the online cost of computing $\mathbf{K}\mathbf{x}$. If efficient online operation is the major concern, evaluating $\mathbf{K}\mathbf{x}$ is fastest when adopting a selection matrix $\mathbf{K} := \mathbf{C} \in \mathcal{C}_{pn}$. However, according to these traditional sketching techniques, selecting samples of $\mathbf{x}$ using $\mathbf{C}$ would imply also a selection (of rows) of $\mathbf{H}$ as $\mathbf{C}\mathbf{H}$, which prevents an optimum (off-line) design of $\mathbf{H}_s$ (cf. the third problem outlined in Section II-A). Finally, note also that traditionally $\mathbf{K}$ is constructed as a so-termed random projection matrix. This facilitates – by leveraging the celebrated Johnson-Lindenstrauss lemma [28] – providing *probabilistic* approximation guarantees that depend on the number p of samples taken.

On the other hand, (selection) sampling picks a subset of p elements of $\mathbf{x}$ in such a way that, if the graph signal is bandlimited with bandwidth k, then it can be exactly reconstructed from the selected samples —the value of the signal at only $p \geq k$ nodes [11]. In essence, a selection matrix $\mathbf{C} \in \mathcal{C}_{pn}$ is designed in such a way that $\mathbf{C}\mathbf{x}$ contains enough information for accurately recovering the whole signal $\mathbf{x}$ via suitable interpolation of these samples [29]. In the noiseless case, any $\mathbf{C}$ with $p \geq k$ that satisfies that $\mathbf{C}\mathbf{V}_k$ is full rank suffices, and perfect recovery is guaranteed by

$$\mathbf{x} = \mathbf{V}_k(\mathbf{C}\mathbf{V}_k)^\dagger(\mathbf{C}\mathbf{x}). \tag{9}$$

In the noisy case, there are several methods and algorithms for obtaining the optimal selection matrix $\mathbf{C}$, see [11], [12]. Selection sampling of $\mathbf{x}$ is particularly useful for compression and storage in the context of streaming data.

All in all, existing sampling schemes consider the properties of $\mathbf{x}$ with the sole objective of signal reconstruction. Accordingly, they are agnostic to higher-level tasks for which the sampled signal $\mathbf{C}\mathbf{x}$ may serve as input. On the other hand, traditional sketching methods neither consider the properties of $\mathbf{x}$, nor they anticipate differences in online and off-line computational power (i.e., focus is on batch processing). Hence, they yield suboptimal alternatives when seeking for an approximate solution with low computational cost in the online phase. This serves as motivation for the approach proposed in this paper, which takes into account the information of $\mathbf{x}$ and $\mathbf{H}$ jointly.

III. Direct sketching as graph signal sampling

After introducing notation and presenting the problem formulation (cf. Fig. 1 and Section II-A1), the focus in this section is on describing the optimal sampling and sketching solutions for the direct problem.

We start by analyzing the solution for the case where the observations are noise free (i.e., $\mathbf{w} = \mathbf{0}$ in Fig. 1). This will be useful to gain insights on the solution to the noisy formulation as well as to derive heuristics to approximate the output of the linear transform. Since in this noiseless scenario we have that $\mathbf{w} = \mathbf{0}$, then the desired output is $\mathbf{y} = \mathbf{H}\mathbf{x}$ and the reduced-complexity approximation is given by $\hat{\mathbf{y}} = \mathbf{H}_s\mathbf{C}\mathbf{x}$. The following result asserts that perfect estimation is possible, namely that $\hat{\mathbf{y}} = \mathbf{y}$ if the number of samples is $p \geq k$.

Proposition 1. *Let $\mathbf{x} \in \mathbb{R}^n$ be a k-bandlimited signal with known spectral template $\mathbf{T} \in \mathbb{C}^{k\times k}$ and let $\mathbf{H} \in \mathbb{R}^{m\times n}$ be a linear transformation. Let $\mathbf{H}_s \in \mathbb{R}^{m\times p}$ be a reduced-input dimensionality sketch of $\mathbf{H}$, $p \leq n$ and $\mathbf{C} \in \mathcal{C}_{pn}$ be a selection matrix. In the absence of noise ($\mathbf{w} = \mathbf{0}$), if $p = k$ and $\mathbf{C}^*$ is designed such that* $\text{rank}\{\mathbf{C}^*\mathbf{V}_k\} = p = k$, *then $\hat{\mathbf{y}} = \mathbf{H}_s^*\mathbf{C}^*\mathbf{x}$ is equal to the desired output $\mathbf{y} = \mathbf{H}\mathbf{x}$ provided that the sketch $\mathbf{H}_s^*$ is given by*

$$\mathbf{H}_s^* = \mathbf{H}\mathbf{V}_k(\mathbf{C}^*\mathbf{V}_k)^{-1}. \tag{10}$$

Proof. The mean-squared error (MSE) criterion [cf. (6)] is

$$\begin{aligned}\mathbb{E}\left[\|\mathbf{y} - \hat{\mathbf{y}}\|_2^2\right] &= \mathbb{E}\left[\|\mathbf{H}\mathbf{x} - \mathbf{H}_s\mathbf{C}\mathbf{x}\|_2^2\right] \\ &= \text{tr}\big[\mathbf{H}\mathbf{R}_x\mathbf{H}^T - 2\mathbf{H}_s\mathbf{C}\mathbf{R}_x\mathbf{H}^T + \mathbf{H}_s\mathbf{C}\mathbf{R}_x\mathbf{C}^T\mathbf{H}_s^T\big].\end{aligned} \tag{11}$$

Recall that $\text{rank}\{\mathbf{C}\} = p$ for any $\mathbf{C} \in \mathcal{C}_{pn}$. Optimizing the MSE cost over $\mathbf{H}_s$ first, results in the following identity that $\mathbf{H}_s^*$ should satisfy (recall $\mathbf{R}_x = \mathbf{V}_k\mathbf{T}\mathbf{V}_k^H$)

$$\mathbf{H}\mathbf{V}_k\mathbf{T}\mathbf{V}_k^H\mathbf{C}^T = \mathbf{H}_s^*\mathbf{C}\mathbf{V}_k\mathbf{T}\mathbf{V}_k^H\mathbf{C}^T. \tag{12}$$

Observe that if $\mathbf{C}$ is chosen such that $\text{rank}\{\mathbf{C}\mathbf{V}_k\} = k$ (i.e. that the p selected rows of $\mathbf{V}_k$ contain a basis for $\mathbb{C}^k$), then $(\mathbf{V}_k^H\mathbf{C}^T\mathbf{C}\mathbf{V}_k)^{-1}$ exists. Now, consider the $p \times k$ matrix $\mathbf{M} = \mathbf{C}\mathbf{V}_k(\mathbf{V}_k^H\mathbf{C}^T\mathbf{C}\mathbf{V}_k)^{-1}\mathbf{T}^{-1}$. If we set $p = k$, then $\mathbf{M}$ is nonsingular and we can multiply both sides of (12) from the right with $\mathbf{M}$, keeping the equivalence between equations. That multiplication yields

$$\mathbf{H}\mathbf{V}_k = \mathbf{H}_s^*\mathbf{C}\mathbf{V}_k.$$

Finally, the fact that $\text{rank}\{\mathbf{C}\mathbf{V}_k\} = p = k$ guarantees that $\mathbf{C}\mathbf{V}_k$ is square and nonsingular, so that its inverse exists and can be used to obtain the solution for $\mathbf{H}_s^*$ given in (10). □

All in all, in the absence of noise it suffices to first set $p = k$ to find a selection matrix $\mathbf{C} \in \mathcal{C}_{pn}$ such that $\text{rank}\{\mathbf{C}\mathbf{V}_k\} = p = k$, and then obtain $\mathbf{H}_s$ as per Proposition 1. This ensures that $\mathbf{y}$ can be formed error-free using p samples of $\mathbf{x}$ via $\mathbf{y} = \hat{\mathbf{y}} = \mathbf{H}_s\mathbf{C}\mathbf{x}$. Clearly, selecting $p \geq k$ will also do the job, provided that $\text{rank}\{\mathbf{C}\mathbf{V}_k\} \geq k$. Before moving on to the noisy setting, a remark is in order.

Remark 2 (Signal recovery via selection sampling). Note that $\text{rank}\{\mathbf{C}\mathbf{V}_k\} = k$ is the same condition for exact signal recovery as with selection sampling [cf. (9)]. Hence, the same techniques proposed in [11] for finding a subset of p rows of $\mathbf{V}_k$ that are linearly independent can be used here. This coincidence should not come as a surprise, since in the absence of noise the design of $\mathbf{C}$ *decouple*s from that of $\mathbf{H}_s$. As a result, existing methods to determine the most informative nodes in sampling scenarios are also applicable here [15].

A. Noisy case

Now consider that the noise vector signal $\mathbf{w} \in \mathbb{R}^n$ is random and independent of $\mathbf{x}$, with $\mathbb{E}[\mathbf{w}] = \mathbf{0}$, $\mathbf{R}_w = \mathbb{E}[\mathbf{w}\mathbf{w}^T] \in \mathbb{R}^{n\times n}$ and $\mathbf{R}_w \succ \mathbf{0}$. In this case, we have $\mathbf{y} = \mathbf{H}\mathbf{x}$ and $\hat{\mathbf{y}} = \mathbf{H}_s\mathbf{C}(\mathbf{x} + \mathbf{w})$ (see Fig. 1). As a result, the design of $\mathbf{H}_s$ and $\mathbf{C}$ to minimize (6) must account for the noise model as well. The joint design of $\mathbf{H}_s$ and $\mathbf{C}$ will be addressed as a two-stage optimization that proceeds into three steps: First, the optimal sketch $\mathbf{H}_s$ is written as a *function* of $\mathbf{C}$. Second, such a function is substituted into the joint optimization to yield a problem that depends only on $\mathbf{C}$. Third, the optimal $\mathbf{H}_s^*$ is found using the function in step one and the optimal value of $\mathbf{C}^*$ found in the step two. Defining the covariance matrix of the output $\mathbf{y}$ as $\mathbf{R}_y = \mathbb{E}[\mathbf{y}\mathbf{y}^T] = \mathbf{H}\mathbf{R}_x\mathbf{H}^T \in \mathbb{R}^{m\times m}$, the outcome of this process is described in the following proposition.

Proposition 2. *The solution of the direct linear sketching problem in the presence of noise* (6) *is given by $\mathbf{C}^*$ and $\mathbf{H}_s^* = \mathbf{H}_s^*(\mathbf{C}^*)$, where*

$$\mathbf{H}_s^*(\mathbf{C}) = \mathbf{H}\mathbf{R}_x\mathbf{C}^T\left(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\right)^{-1} \tag{13}$$

and $\mathbf{C}^*$ can be obtained as the solution to the problem

$$\min_{\mathbf{C}\in\mathcal{C}_{pn}} \ \text{tr}\left[\mathbf{R}_y - \mathbf{H}\mathbf{R}_x\mathbf{C}^T\left(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\right)^{-1}\mathbf{C}\mathbf{R}_x\mathbf{H}^T\right]. \tag{14}$$

Proof. The objective function in (6) can be rewritten as

$$\begin{aligned}\mathbb{E}\left[\|\mathbf{y} - \hat{\mathbf{y}}\|_2^2\right] &= \mathbb{E}\left[\|\mathbf{H}\mathbf{x} - \mathbf{H}_s\mathbf{C}(\mathbf{x} + \mathbf{w})\|_2^2\right] \\ &= \text{tr}\left[\mathbf{H}\mathbf{R}_x\mathbf{H}^T - 2\mathbf{H}_s\mathbf{C}\mathbf{R}_x\mathbf{H}^T + \mathbf{H}_s\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\mathbf{H}_s^T\right]\end{aligned} \tag{15}$$

since $\mathbf{x}$ and $\mathbf{w}$ are assumed independent. Optimizing with respect to $\mathbf{H}_s$ first, yields

$$\mathbf{H}_s^*(\mathbf{C}) = \mathbf{H}\mathbf{R}_x\mathbf{C}^T\left(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\right)^{-1}$$

establishing (13). Matrix $\mathbf{C} \in \mathcal{C}_{pn}$ is full rank since it selects p *distinct* nodes, then $\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T$ has rank p and thus it is invertible [30]. Substituting the expression for $\mathbf{H}_s^*(\mathbf{C})$ into (15), yields (14). □

Indeed, Proposition 2 confirms that the optimal selection matrix $\mathbf{C}^*$ is obtained solving problem (14) which, after leveraging the expression in (13), only requires knowledge of the given matrices $\mathbf{H}$, $\mathbf{R}_x$ and $\mathbf{R}_w$. The optimal sketch is then found substituting the obtained $\mathbf{C}^*$ into (13) as $\mathbf{H}_s^* = \mathbf{H}_s^*(\mathbf{C}^*)$.

Furthermore, the next proposition shows that problem (14), which yields the optimal $\mathbf{C}^*$, is equivalent to a *binary* optimization problem, with a linear objective function and subject to linear matrix inequality (LMI) constraints.

Proposition 3. *Let $\mathbf{c} \in \{0,1\}^n$ be a binary vector such that $\mathbf{C}^T\mathbf{C} = \text{diag}(\mathbf{c})$. Then, in the context of Proposition 2, the optimization problem* (14) *over $\mathbf{C}$ is equivalent to*

$$\begin{aligned}\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}_\alpha}} \ & \text{tr}\left[\mathbf{Y}\right] \\ \text{s.t. } & \bar{\mathbf{C}}_\alpha = \alpha^{-1}\text{diag}(\mathbf{c}) \ , \ \mathbf{c}^T\mathbf{1}_n = p \\ & \begin{bmatrix}\mathbf{Y} - \mathbf{R}_y + \mathbf{H}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{H}^T & \mathbf{H}\mathbf{R}_x\bar{\mathbf{C}}_\alpha \\ \bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{H}^T & \bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha\end{bmatrix} \succeq \mathbf{0}\end{aligned} \tag{16}$$

where $\bar{\mathbf{R}}_\alpha = (\mathbf{R}_x + \mathbf{R}_w - \alpha \mathbf{I}_n)$, $\bar{\mathbf{C}}_\alpha$ *and* $\mathbf{Y} \in \mathbb{R}^{m\times m}$ *are auxiliary variables and* $\alpha > 0$ *is any scalar satisfying* $\bar{\mathbf{R}}_\alpha \succ \mathbf{0}$.

Proof. The inverse in the objective of (14) can be written as [15]

$$(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T)^{-1} = \left(\alpha \mathbf{I}_p - \alpha \mathbf{I}_p + \mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\right)^{-1}$$
$$= \alpha^{-1}\mathbf{I}_p - \alpha^{-2}\mathbf{C}(\bar{\mathbf{R}}_\alpha^{-1} + \alpha^{-1}\mathbf{C}^T\mathbf{C})^{-1}\mathbf{C}^T, \quad (17)$$

where $\alpha \neq 0$ is a rescaling parameter, and in obtaining the second equality we used the Woodbury Matrix Identity [27]. Note that α has to be such that $\bar{\mathbf{R}}_\alpha = (\mathbf{R}_x + \mathbf{R}_w - \alpha \mathbf{I}_n)$ is still invertible. Substituting (17) into (14) and recalling that $\mathbf{C}^T\mathbf{C} = \text{diag}(\mathbf{c})$, we have that

$$\min_{\mathbf{c}\in\{0,1\}^n, \bar{\mathbf{C}}_\alpha} \text{tr}\Big[\mathbf{R}_y - \mathbf{H}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{H}^T \quad (18)$$
$$+ \mathbf{H}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\left(\bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha\right)^{-1}\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{H}^T\Big]$$
$$\text{s.t. } \bar{\mathbf{C}}_\alpha = \alpha^{-1}\text{diag}(\mathbf{c}) \ , \ \mathbf{c}^T\mathbf{1}_n = p.$$

Note that in (18) we optimize over a binary vector $\mathbf{c} \in \mathbb{R}^n$ with exactly p nonzero entries, instead of a binary matrix $\mathbf{C} \in \mathcal{C}_{pn}$. The p nonzero elements in $\mathbf{c}$ indicate the nodes to be sampled. Problem (18) can be reformulated as

$$\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}_\alpha}} \text{tr}\left[\mathbf{Y}\right] \quad (19)$$
$$\text{s.t. } \mathbf{R}_y - \mathbf{H}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{H}^T$$
$$+ \mathbf{H}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\left(\bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha\right)^{-1}\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{H}^T \preceq \mathbf{Y}$$
$$\bar{\mathbf{C}}_\alpha = \alpha^{-1}\text{diag}(\mathbf{c}) \ , \ \mathbf{c}^T\mathbf{1}_n = p$$

where $\mathbf{Y} \in \mathbb{R}^{m\times m}$, $\mathbf{Y} \succeq 0$ is an auxiliary optimization variable. Using the Schur-complement lemma for *positive definite* matrices [31], problem (19) can be written as (16). Hence, to complete the proof we need to show that $\bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha \succeq \mathbf{0}$ so that the aforementioned lemma can be invoked. To that end, suppose first that $\alpha < 0$. Then we have that $\bar{\mathbf{R}}_\alpha \succ \mathbf{0}$ and $\bar{\mathbf{C}}_\alpha = \alpha^{-1}\text{diag}(\mathbf{c}) \preceq \mathbf{0}$, so that $\bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha \succeq \mathbf{0}$ may not be positive definite. Suppose now that $\alpha > 0$. Then $\bar{\mathbf{C}}_\alpha \succeq \mathbf{0}$ and there always exists a sufficiently small positive α such that $\bar{\mathbf{R}}_\alpha \succ \mathbf{0}$ since $\mathbf{R}_w \succ \mathbf{0}$. This implies that if α is chosen such that $\alpha > 0$ and $\bar{\mathbf{R}}_\alpha = \mathbf{R}_x + \mathbf{R}_w - \alpha\mathbf{I}_n \succ \mathbf{0}$ (which are the conditions stated in the proposition), then $\bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha$ is positive definite and problems (18) and (16) are equivalent. □

Problem (16) is an SDP optimization modulo the binary constraints on $\mathbf{C}$. Note that, from an analytical perspective, the two-stage optimization approach requires first finding the optimal sketch $\mathbf{H}_s^*$ as a function of $\mathbf{C}$ via (13) and then substitute $\mathbf{H}_s^*(\mathbf{C})$ into (6) to yield the optimization in (14) or (16). From an algorithmic perspective, the order is however reversed. First, we find $\mathbf{C}$ by "solving" the binary optimization in either (14) or (16), and then the resulting $\mathbf{C}^*$ is substituted into (13) to find $\mathbf{H}_s^*$ in $O(mnp + p^3)$ complexity. A convex relaxation of (16) as well as other heuristic approaches to obtain a solution are presented in Section V.

B. Sampling the linear transform

As explained in Section II-A3, there may be setups where designing the new operator $\mathbf{H}_s$ in (13), whose entries do not resemble those in $\mathbf{H}$, is costly or even unfeasible. An alternative to bypass this problem consists in forming the sketch by taking p *columns* of $\mathbf{H}$, i.e., setting $\mathbf{H}_s = \mathbf{H}\mathbf{C}^T$. Although from an MSE performance point of view such a *sketching* design is suboptimal [cf. (13)], numerical tests carried out in Section VI suggest its can sometimes yield competitive performance. The optimal *sampling* strategy for sketches within this class is given in the following proposition.

Proposition 4. *Let* $\mathbf{H}_s = \mathbf{H}\mathbf{C}^T$ *be constructed from a subset of* p *columns of* $\mathbf{H}$. *Then, the optimal sampling matrix* $\mathbf{C}^*$ *can be recovered from* $\mathbf{c}^*$ *such that* $\text{diag}(\mathbf{c}^*) = (\mathbf{C}^*)^T\mathbf{C}^*$, *where* $\mathbf{c}^*$ *is the solution to the following problem*

$$\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}}} \text{tr}[\mathbf{Y}] \quad (20)$$
$$\text{s.t. } \bar{\mathbf{C}} = \text{diag}(\mathbf{c}) \ , \ \mathbf{c}^T\mathbf{1}_n = p$$
$$\begin{bmatrix} \mathbf{Y} - \mathbf{R}_y + 2\mathbf{H}\bar{\mathbf{C}}\mathbf{R}_x\mathbf{H}^T & \mathbf{H}\bar{\mathbf{C}} \\ \bar{\mathbf{C}}\mathbf{H}^T & (\mathbf{R}_x + \mathbf{R}_w)^{-1} \end{bmatrix} \succeq \mathbf{0}.$$

Proof. Defining the matrix $\bar{\mathbf{C}} = \text{diag}(\mathbf{c})$, the estimated output is

$$\hat{\mathbf{y}} = \mathbf{H}\mathbf{C}^T\mathbf{C}(\mathbf{x} + \mathbf{w}) = \mathbf{H}\text{diag}(\mathbf{c})(\mathbf{x} + \mathbf{w}) = \mathbf{H}\bar{\mathbf{C}}(\mathbf{x} + \mathbf{w}).$$

Since the desired output is $\mathbf{y} = \mathbf{H}\mathbf{x}$, the MSE objective function simplifies to

$$\mathbb{E}\left[\|\mathbf{H}\mathbf{x} - \mathbf{H}\bar{\mathbf{C}}(\mathbf{x} + \mathbf{w})\|_2^2\right] \quad (21)$$
$$= \text{tr}\left[\mathbf{H}\mathbf{R}_x\mathbf{H}^T - 2\mathbf{H}\bar{\mathbf{C}}\mathbf{R}_x\mathbf{H}^T + \mathbf{H}\bar{\mathbf{C}}(\mathbf{R}_x + \mathbf{R}_w)\bar{\mathbf{C}}\mathbf{H}^T\right].$$

Introducing the auxiliary variable $\mathbf{Y} \in \mathbb{R}^{m\times m}$, $\mathbf{Y} \succeq \mathbf{0}$, minimizing (21) with $\mathbf{C} \in \mathcal{C}_{pn}$ is equivalent to solving

$$\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}}} \text{tr}\left[\mathbf{Y}\right] \quad (22)$$
$$\text{s.t. } \bar{\mathbf{C}} = \text{diag}(\mathbf{c}) \ , \ \mathbf{c}^T\mathbf{1}_n = p$$
$$\mathbf{R}_y - 2\mathbf{H}\bar{\mathbf{C}}\mathbf{R}_x\mathbf{H}^T + \mathbf{H}\bar{\mathbf{C}}(\mathbf{R}_x + \mathbf{R}_w)\bar{\mathbf{C}}\mathbf{H}^T \preceq \mathbf{Y}.$$

Finally, using the Schur complement-based lemma for positive semidefiniteness [31], then (22) can be shown equivalent to (20), competing the proof. □

IV. Inverse sketching problem

This section presents the optimal sampling $\mathbf{C}^* \in \mathcal{C}_{np}$ and sketching $\mathbf{H}_s^* \in \mathbb{R}^{m\times p}$ solutions for the inverse problem in Fig. 2 and Section II-A2. Paralleling the structure of the previous section, we first analyze the particular case where the observations are noise free, then shift to the general noisy case, and finally look at a suboptimal scheme where the sketch is formed by taking *rows* of the linear transformation $\mathbf{H}$.

If there is no noise corrupting the output observations, then (8) simplifies to

$$\min_{\mathbf{C}\in\mathcal{C}_{np}, \mathbf{H}_s} \mathbb{E}\left[\|\mathbf{H}(\mathbf{H}_s\mathbf{C}\mathbf{x}) - \mathbf{x}\|_2^2\right]. \quad (23)$$

Assuming that $\mathbf{H}$ is full rank – so that matrix $\mathbf{A}_{\text{LS}} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T$ in (7) exists –, the optimal $\mathbf{C}^*$ and $\mathbf{H}_s^*$ solving

(23) are characterized in the following proposition, which shows that once more perfect reconstruction is possible.

Proposition 5. *Let $\mathbf{x} \in \mathbb{R}^n$ be a k-bandlimited signal with known spectral template $\mathbf{T} \in \mathbb{C}^{k\times k}$ and let $\mathbf{H} \in \mathbb{R}^{n\times m}$ be a linear transformation. Let $\mathbf{H}_s \in \mathbb{R}^{m\times p}$ be a reduced-input dimensionality sketch, $p \leq n$ and $\mathbf{C} \in \mathcal{C}_{pn}$ be a selection matrix. If $p = k$, a sampling $\mathbf{C}^*$ designed such that* $\text{rank}\{\mathbf{C}^*\mathbf{V}_k\} = p = k$ *is optimal for* (23)*. Moreover, $\hat{\mathbf{y}} = \mathbf{H}_s^*\mathbf{C}^*\mathbf{x}$ yields the LS estimate $\hat{\mathbf{y}}_{\text{LS}} = \mathbf{A}_{\text{LS}}\mathbf{x} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{x}$ provided that the optimum $\mathbf{H}_s^*$ is given by*

$$\mathbf{H}_s^*(\mathbf{C}^*) = \mathbf{A}_{\text{LS}}\mathbf{V}_k(\mathbf{C}^*\mathbf{V}_k)^{-1}. \tag{24}$$

Proof. The MSE criterion [cf. (23)] is

$$\begin{aligned}&\mathbb{E}\left[\|\mathbf{H}(\mathbf{H}_s\mathbf{C}\mathbf{x}) - \mathbf{x}\|_2^2\right] \\ &= \text{tr}\left[\mathbf{H}\mathbf{H}_s\mathbf{C}\mathbf{R}_x\mathbf{C}^T\mathbf{H}_s^T\mathbf{H}^T - 2\mathbf{H}\mathbf{H}_s\mathbf{C}\mathbf{R}_x + \mathbf{R}_x\right].\end{aligned} \tag{25}$$

Observe that because $\mathbf{C} \in \mathcal{C}_{pn}$ then $\text{rank}\{\mathbf{C}\} = p$. Now, optimizing (25) over $\mathbf{H}_s$ first and recalling that $\mathbf{A}_{\text{LS}} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T$ and $\mathbf{R}_x = \mathbf{V}_k\mathbf{T}\mathbf{V}_k^H$, results in

$$\mathbf{H}_s^*\mathbf{C}\mathbf{V}_k\mathbf{T}\mathbf{V}_k^H\mathbf{C}^T = \mathbf{A}_{\text{LS}}\mathbf{V}_k\mathbf{T}\mathbf{V}_k^H\mathbf{C}^T. \tag{26}$$

Now, choose $\mathbf{C}$ such that $\text{rank}\{\mathbf{C}\mathbf{V}_k\} = k$, so that $(\mathbf{V}_k^H\mathbf{C}^T\mathbf{C}\mathbf{V}_k)$ is a nonsingular $k \times k$ matrix. Then, by setting $p = k$, we have that the $k \times k$ matrix $\mathbf{M} = \mathbf{C}\mathbf{V}_k(\mathbf{V}_k^H\mathbf{C}^T\mathbf{C}\mathbf{V}_k)^{-1}\mathbf{T}^{-1}$ is nonsingular, so that by multiplying both sides of (26) from the right with $\mathbf{M}$, one arrives at

$$\mathbf{H}_s^*\mathbf{C}\mathbf{V}_k = \mathbf{A}_{\text{LS}}\mathbf{V}_k.$$

Finally, because of the choice $p = k$ and the design of $\mathbf{C}$ such that $\text{rank}\{\mathbf{C}\mathbf{V}_k\} = p = k$, then $(\mathbf{C}\mathbf{V}_k)$ is $k \times k$ matrix with an inverse. Consequently, the closed form solution for $\mathbf{H}_s^*$ is that given by (24). □

Note that in this noise-free case, the result of Proposition 5 is identical to that of Proposition 1 with a direct filter $\mathbf{H}_{\text{F}}$ that is equal to $\mathbf{H}_{\text{F}} = \mathbf{A}_{\text{LS}}$. This is consistent with the fact that in the present setting, $\mathbf{H}_s$ is a sketch of the inverse mapping relating outputs to inputs.

A. Noisy case

As in Section III-A we consider that the noise $\mathbf{w} \in \mathbb{R}^n$ is independent of $\mathbf{x}$, has zero mean and that its covariance matrix satisfies $\mathbf{R}_w \succ \mathbf{0}$. Upon defining $\mathbf{G} = \mathbf{H}\mathbf{A}_{\text{LS}}$, the next proposition characterizes the optimal $\mathbf{C}^*$ and $\mathbf{H}_s^*$ as a two-step process.

Proposition 6. *If $\mathbf{H}$ is full rank, the solution to problem* (8) *is given by $\mathbf{C}^*$ and $\mathbf{H}_s^* = \mathbf{H}_s^*(\mathbf{C}^*)$, where*

$$\mathbf{H}_s^*(\mathbf{C}) = \mathbf{A}_{\text{LS}}\mathbf{R}_x\mathbf{C}^T\left(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\right)^{-1}. \tag{27}$$

The optimum selection matrix $\mathbf{C}^*$ is obtained from solving

$$\min_{\mathbf{C}\in\mathcal{C}_{np}} \text{tr}\big[\mathbf{R}_x - \mathbf{G}\mathbf{R}_x\mathbf{C}^T\left(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\right)^{-1}\mathbf{C}\mathbf{R}_x\mathbf{G}^T\big]. \tag{28}$$

Proof. In the noisy setting, the MSE cost can be expanded as

$$\begin{aligned}&\mathbb{E}\left[\left\|\mathbf{H}\mathbf{H}_s\mathbf{C}(\mathbf{x} + \mathbf{w}) - \mathbf{x}\right\|_2^2\right] \\ &= \text{tr}\big[\mathbf{R}_x - 2\mathbf{H}\mathbf{H}_s\mathbf{C}\mathbf{R}_x + \mathbf{H}\mathbf{H}_s\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\mathbf{H}_s^T\mathbf{H}^T\big].\end{aligned} \tag{29}$$

Taking the derivative of (29) with respect to $\mathbf{H}_s$ and setting it equal to zero yields the optimal sketch

$$\mathbf{H}_s^*(\mathbf{C}) = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{R}_x\mathbf{C}^T\left(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T\right)^{-1}.$$

Substituting the definition of $\mathbf{A}_{\text{LS}}$ into the previous yields equation (27). Substituting now (27) into (29) simplifies the optimization problem (8) into (28). □

As was the case for the direct problem, the joint design of the optimal selection matrix $\mathbf{C}^*$ and the optimal sketch $\mathbf{H}_s^*$ can be carried out as a two-step process. First $\mathbf{C}^*$ is found as the solution to (28), which only requires as input $\mathbf{R}_x$, $\mathbf{G}$ and $\mathbf{R}_w$. Then, the optimal sketch is found using (27) $\mathbf{H}_s^* = \mathbf{H}_s^*(\mathbf{C}^*)$.

Moreover, note that the optimal sketch $\mathbf{H}_s^*$ is tantamount to the LS operator $\mathbf{A}_{\text{LS}} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T$ acting on a preprocessed version of $\mathbf{x}$, via the matrix $\mathbf{R}_x\mathbf{C}^T(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T)^{-1}$. What this preprocessing entails is, essentially, choosing the samples of $\mathbf{x}$ with the optimal tradeoff between the signal in the multiplicative term $\mathbf{R}_x\mathbf{C}^T$ and the noise in the inverse term $(\mathbf{C}(\mathbf{R}_x + \mathbf{R}_w)\mathbf{C}^T)^{-1}$.

Also observe that, while computation of the inverse-mapping sketch $\mathbf{H}_s^*$ involves the (costly) computation of $(\mathbf{H}^T\mathbf{H})^{-1}$, this is carried out entirely off line.

Finally, the optimization problem (28) over $\mathbf{C}$ is equivalent to a problem reminiscent of an SDP save the binary constraints, as shown next.

Proposition 7. *In the setting of Proposition 6, the optimum selection matrix $\mathbf{C}^*$ is such that $(\mathbf{C}^*)^T\mathbf{C}^* = \text{diag}(\mathbf{c}^*)$, with $\mathbf{c}^*$ being the the solution to the problem*

$$\begin{aligned}\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}_\alpha}}\ &\text{tr}\left[\mathbf{Y}\right] \\ \text{s.t. } &\bar{\mathbf{C}}_\alpha = \alpha^{-1}\text{diag}(\mathbf{c})\ ,\ \mathbf{c}^T\mathbf{1}_n = p \\ &\begin{bmatrix}\mathbf{Y} - \mathbf{R}_x + \mathbf{G}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{G}^T & \mathbf{G}\mathbf{R}_x\bar{\mathbf{C}}_\alpha \\ \bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{G}^T & \bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha\end{bmatrix} \succeq \mathbf{0},\end{aligned} \tag{30}$$

where $\bar{\mathbf{R}}_\alpha = (\mathbf{R}_x + \mathbf{R}_w - \alpha\mathbf{I}_n)$, $\bar{\mathbf{C}}_\alpha$ and $\mathbf{Y} \in \mathbb{R}^{n\times n}$ are auxiliary variables and $\alpha > 0$ is any scalar satisfying $\bar{\mathbf{R}}_\alpha \succ \mathbf{0}$.

Proof. Consider the nonzero parameter $\alpha \in \mathbb{R}$. Then, by using (17) into (28) we get

$$\begin{aligned}\min_{\mathbf{c}\in\{0,1\}^n,\bar{\mathbf{C}}_\alpha}\ &\text{tr}\Big[\mathbf{R}_x - \mathbf{G}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{G}^T \\ &\qquad + \mathbf{G}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\left(\bar{\mathbf{R}}_\alpha^{-1} + \bar{\mathbf{C}}_\alpha\right)^{-1}\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{G}^T\Big] \\ \text{s.t. } &\bar{\mathbf{C}}_\alpha = \alpha^{-1}\text{diag}(\mathbf{c})\ ,\ \mathbf{c}^T\mathbf{1}_n = p.\end{aligned} \tag{31}$$

Let $\mathbf{Y} \in \mathbb{R}^{n\times n}$ be a new variable such that $\mathbf{Y} \succeq \mathbf{0}$. Then (31) is equivalent to

$$\begin{aligned}\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}_\alpha}} \ & \mathrm{tr}\left[\mathbf{Y}\right] \\ \text{s.t. } & \mathbf{R}_x - \mathbf{G}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{G}^T \\ & \quad + \mathbf{G}\mathbf{R}_x\bar{\mathbf{C}}_\alpha\left(\bar{\mathbf{R}}_\alpha^{-1}+\bar{\mathbf{C}}_\alpha\right)^{-1}\bar{\mathbf{C}}_\alpha\mathbf{R}_x\mathbf{G}^T \preceq \mathbf{Y} \\ & \bar{\mathbf{C}}_\alpha = \alpha^{-1}\mathrm{diag}(\mathbf{c})\ ,\ \mathbf{c}^T\mathbf{1}_n = p.\end{aligned} \tag{32}$$

Using the Schur complement lemma for semidefinite matrices, (32) is finally equivalent to (30). As was the case for the direct problem, the fixed parameter $\alpha > 0$ has to be chosen so that $(\bar{\mathbf{R}}_\alpha + \bar{\mathbf{C}}_\alpha) \succeq \mathbf{0}$ for all $\mathbf{c}$. Otherwise, the feasible set of problem (30) will be reduced to those $\mathbf{c}$ such that $(\bar{\mathbf{R}}_\alpha + \bar{\mathbf{C}}_\alpha) \succeq \mathbf{0}$ breaking the equivalence between (30) and (32). □

B. Sampling the linear transformation

Here we look at setups where the sketch $\mathbf{H}_s$ is constrained to be a submatrix of $\mathbf{H}$. In the context of linear inverse problems, this amounts to saying that $\mathbf{H}_s = (\mathbf{C}\mathbf{H})^T$ is formed using p rows of $\mathbf{H}$. With $\bar{\mathbf{H}} = \mathbf{H}\mathbf{H}^T$ denoting the $n\times n$ Gram matrix of $\mathbf{H}^T$, the solution to (8) for sketches of the form $\mathbf{H}_s = (\mathbf{C}\mathbf{H})^T$ is given by the following proposition.

Proposition 8. *Let $\mathbf{H}_s = (\mathbf{C}\mathbf{H})^T$ be constructed from a subset of p rows of $\mathbf{H}$. Then, the optimal sampling matrix $\mathbf{C}^*$ can be recovered from $\mathbf{c}^*$ such that $\mathrm{diag}(\mathbf{c}^*) = (\mathbf{C}^*)^T\mathbf{C}^*$, where $\mathbf{c}^*$ is the solution to the following problem*

$$\begin{aligned}\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}}} \ & \mathrm{tr}[\mathbf{Y}] \\ \text{s.t. } & \bar{\mathbf{C}} = \mathrm{diag}(\mathbf{c})\ ,\ \mathbf{c}^T\mathbf{1}_n = p \\ & \begin{bmatrix}\mathbf{Y} - \mathbf{R}_x + 2\bar{\mathbf{H}}\bar{\mathbf{C}}\mathbf{R}_x & \bar{\mathbf{H}}\bar{\mathbf{C}} \\ \bar{\mathbf{C}}\bar{\mathbf{H}} & (\mathbf{R}_x+\mathbf{R}_w)^{-1}\end{bmatrix} \succeq \mathbf{0}.\end{aligned} \tag{33}$$

Proof. With $\bar{\mathbf{C}} = \mathrm{diag}(\mathbf{c})$ the estimated response is $\mathbf{H}(\mathbf{C}\mathbf{H})^T\mathbf{C}(\mathbf{x}+\mathbf{w}) = \bar{\mathbf{H}}\bar{\mathbf{C}}(\mathbf{x}+\mathbf{w})$. Given that the desired response is $\mathbf{x} = \mathbf{H}\mathbf{y}$, the MSE cost function simplifies to

$$\begin{aligned}&\mathbb{E}\left[\|\mathbf{x} - \bar{\mathbf{H}}\bar{\mathbf{C}}(\mathbf{x}+\mathbf{w})\|_2^2\right] \\ &= \mathrm{tr}\left[\mathbf{R}_x - 2\bar{\mathbf{H}}\bar{\mathbf{C}}\mathbf{R}_x + \bar{\mathbf{H}}\bar{\mathbf{C}}(\mathbf{R}_x+\mathbf{R}_w)\bar{\mathbf{C}}\bar{\mathbf{H}}\right]\end{aligned} \tag{34}$$

Upon introducing the auxiliary optimization variable $\mathbf{Y} \in \mathbb{R}^{n\times n}$, $\mathbf{Y} \succeq \mathbf{0}$, minimizing (34) with $\mathbf{C} \in \mathcal{C}_{pn}$ can be equivalently reformulated as

$$\begin{aligned}\min_{\substack{\mathbf{c}\in\{0,1\}^n,\\ \mathbf{Y},\bar{\mathbf{C}}}} \ & \mathrm{tr}\left[\mathbf{Y}\right] \\ \text{s.t. } & \bar{\mathbf{C}} = \mathrm{diag}(\mathbf{c})\ ,\ \mathbf{c}^T\mathbf{1}_n = p \\ & \mathbf{R}_x - 2\bar{\mathbf{H}}\bar{\mathbf{C}}\mathbf{R}_x + \bar{\mathbf{H}}\bar{\mathbf{C}}(\mathbf{R}_x+\mathbf{R}_w)\bar{\mathbf{C}}\bar{\mathbf{H}} \preceq \mathbf{Y}.\end{aligned} \tag{35}$$

Once more, based on Schur complement arguments one can readily establish that (35) is equivalent to (33). □

Method	Number of operations	Time [s]
Optimal solution	$O\left(\binom{n}{p}\right)$	$> 10^6$
SDP relaxation	$O\left((m+n)^{3.5}\right)$	25.29
Noise-blind heuristic	$O\left(n\log n\right)$	0.46
Greedy approach	$O\left(mn^2p^2 + np^4\right)$	174.38

Table I: Computational complexity of the heuristic methods proposed in Section V to solve optimization problems in Sections III and IV: (i) Number of operations required (ii) Time (in seconds) required to solve the problem in Section VI-C, see Fig 5.

V. Heuristic approaches

In this section, several heuristics are outlined for tackling the linear sketching problems described so far. The rationale behind this is that oftentimes the problems posed in Sections III and IV can be intractable, or, just too computationally expensive even if carried out off line. In fact, the optimal solution $\mathbf{C}^*$ to (14) or (28) can be obtained by evaluating the objective function in each one of the $\binom{n}{p}$ possible solutions. Table I lists the complexity of each of the proposed methods. Additionally, the time (in seconds) taken to run the simulation related to Fig. 5 is also included in the table for comparison. In all cases, after obtaining $\mathbf{C}^*$, the optimal value of $\mathbf{H}_s^*$ entails $O(mnp + p^3)$ operations, see (13) and (27).

A. Convex relaxation

Recall that the main difficulty when solving (16), (20), (30) and (33) are the *binary* constraints that render the problems non-convex and, in fact, NP-hard. A standard alternative to overcome this difficulty is to relax the binary constraint on the sampling vector $\mathbf{c} \in \{0,1\}^n$ as

$$\mathbf{c} \in [0,1]^n. \tag{36}$$

This way, the optimization problems become convex and can be solved with polynomial complexity in $O((m+n)^{3.5})$ operations as per the SDP formulations in problems (16) and (30) [31].

Once a solution to the relaxed problem is obtained, two ways of recovering a binary vector $\mathbf{c}$ are considered. The first one is a deterministic method referred to as thresholding, which simply consists in setting the largest p elements to 1 and the rest to 0. The second one consists in computing $\mathbf{p}_c = \mathbf{c}/\|\mathbf{c}\|_1$, which can be viewed as a probability distribution over the nodes. The sampled nodes are then drawn at random from this distribution, see [32]. This should be done once, off line, and the same selection matrix used for every incoming input (or output). Although not pursued here, also pertinent are formulations that penalize the objective with $\|\mathbf{c}\|_1$ and leverage ℓ_1-norm minimization advances to promote sparsity on $\mathbf{c}$.

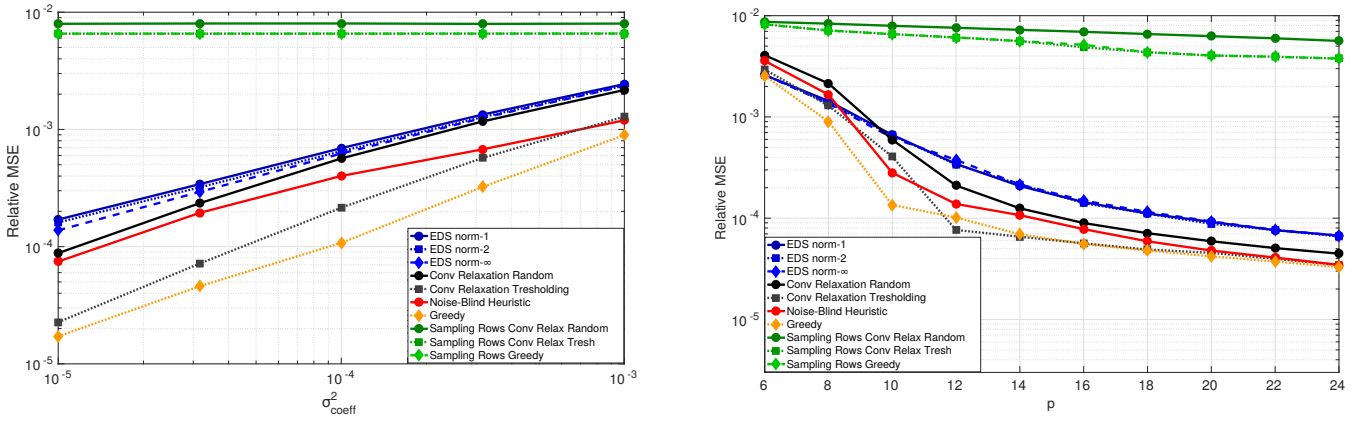

(a) Function of the noise variance

(b) Function of the sample size

Figure 3. Approximating the GFT as an inverse sketching problem. Relative estimated MSE $\|\hat{\mathbf{y}}-\mathbf{y}\|^2/\|\mathbf{y}\|^2$ for the problem of estimating the k frequency components of a bandlimited graph signal on an Erdős-Renyi graph from noisy observations. (a) As a function of the noise variance for fixed $p=k=10$. (b) As a function of the number of samples for a fixed noise variance given by $\sigma^2_{\text{coeff}}=10^{-4}$.

B. Noise-blind heuristic

Recall the cost function of the direct problem (14), namely

$$\text{tr}\left[\mathbf{R}_y-\mathbf{H}\mathbf{R}_x\mathbf{C}^T\left(\mathbf{C}(\mathbf{R}_x+\mathbf{R}_w)\mathbf{C}^T\right)^{-1}\mathbf{C}\mathbf{R}_x\mathbf{H}^T\right]. \quad (37)$$

The first term $\text{tr}[\mathbf{R}_y]=\text{tr}[\mathbf{H}\mathbf{R}_x\mathbf{H}^T]$ is the covariance matrix of the complete filtered signal, which does not depend on $\mathbf{C}$. The second term accounts for the energy of the filtered signal samples $\mathbf{C}\mathbf{R}_x\mathbf{H}^T$, as well as for the residual noise power in those samples $\mathbf{C}(\mathbf{R}_x+\mathbf{R}_w)\mathbf{C}^T$. Ideally, to minimize the objective function (37), the samples selected should be those that simultaneously maximize the output signal energy while minimizing the influence of the noise (and that is why the noise covariance matrix appears in the inverse).

Now, if we ignore the effect of the noise altogether, then selecting the p rows of $\mathbf{R}_x\mathbf{H}^T$ that have maximum ℓ_2 norm should suffice to maximize the second term. Given that the matrix $\mathbf{R}_x\mathbf{H}^T$ has n rows, selecting the p with highest norm is indicative of those p most informative nodes to sample. This heuristic method for designing the selection matrix is straightforward to implement, entails $O(n\log n)$ operations for the sorting algorithm [33] and is shown in Section VI to yield satisfactory performance, especially if the noise variance is low or the filter has a lot of structure.

For the inverse sketching problem, inspection of (28) suggests that an analogous scheme for designing the selection matrix is to pick the p rows of $(\mathbf{G}\mathbf{R}_x)^T$ with largest ℓ_2 norm.

C. Greedy approach

Another alternative to approximate the solution of (14) and (28), or approximate the solution of minimizing (21) and (34) over $\mathcal{C}_{np}$, is to implement an iterative greedy algorithm that adds nodes to the sampling set incrementally, so that each iteration the node that reduces the error the most is incorporated to the sampling set. Considering problem (14) as an example, first, one-by-one all n nodes are tested and the one that yields the lowest value of the objective function in (14) is added to the sampling set. Then, the sampling set is augmented with one more node by choosing the one that yields the lowest optimal objective among the remaining $n-1$ nodes. The procedure is repeated until p nodes are selected in the sampling set. This way only $n+(n-1)+\cdots+(n-(p-1))<np$ evaluations of the objective function (14) are required. Note that each evaluation of (14) entails $O(mnp+p^3)$ operations, so that the overall cost of the greedy approach is $O(mn^2p^2+np^4)$. Greedy algorithms have well-documented merits for sample selection, even for non-submodular objectives like the one in (14); see [16], [34].

VI. Numerical tests

To demonstrate the effectiveness of the sketching methods developed in this paper, two numerical test cases are considered here. In Sections VI-A and VI-B, the GFT coefficients of a bandlimited graph signal are estimated by sampling and applying a linear transform to the signal in the graph domain. In Section VI-C, we classify two digits from the MNIST Handwritten Digit database [35] by directly procesing selected pixels through a sketch of a linear classifier combined with a Principal Component Analysis (PCA) transform.

In all cases, the sketching-as-sampling technique presented in this paper is compared to traditional sampling techniques in noisy environments. Then, the optimal matrix $\mathbf{H}_2^*(\mathbf{C})$ is obtained as in (13) or (27). The sampling matrix $\mathbf{C}$ is obtained according to the experimental design sampling (EDS) method proposed in [36] where weights $\kappa_j=\sup_{1\leq i\leq n}|\sqrt{n}\mathbf{V}_{j,i}|$ are used to construct a probability distribution over the nodes and then sample with replacement according to this distribution. It is important to observe that κ_j is obtained as the ℓ_∞ norm of the rows of $\mathbf{V}$. In this paper, the ℓ_2 and the ℓ_1 norm are also considered. The use of the ℓ_2 norm corresponds to the methods proposed in [32].

A. Approximating the GFT as an inverse problem

Let $\mathcal{G}$ be an Erdős-Renyi graph with n nodes and where each pair of nodes is connected with probability 0.2, independently across pairs. The graph shift operator $\mathbf{S}=\mathbf{V}\mathbf{\Lambda}\mathbf{V}^H$ is set to

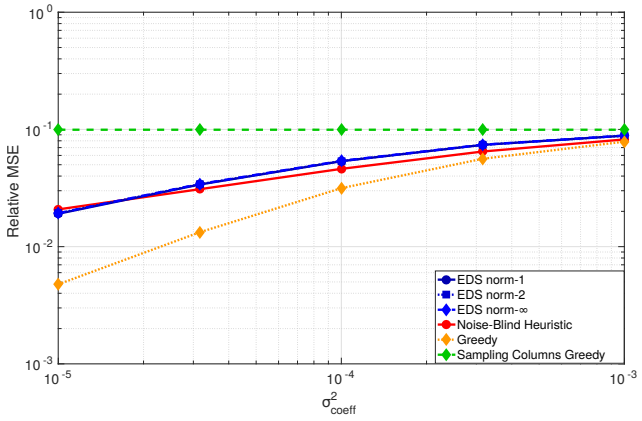

(a) Function of the noise variance

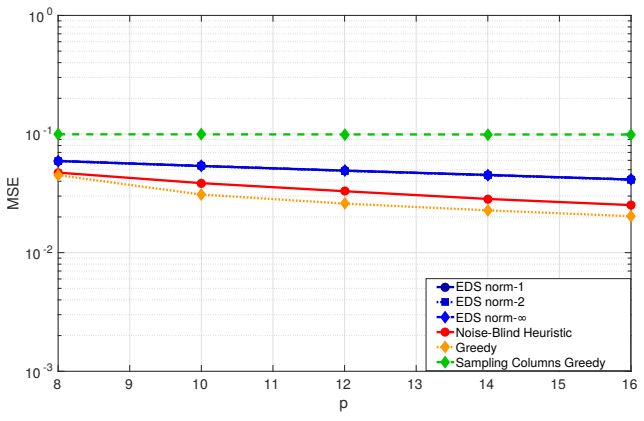

(b) Function of the sample size

Figure 4. Approximating the GFT as a direct sketching problem for a large network. Relative estimated MSE $\|\hat{\mathbf{y}}-\mathbf{y}\|^2/\|\mathbf{y}\|^2$ for the problem of estimating the $k=10$ frequency components of a bandlimited graph signal on an Erdős-Renyi graph with $p=0.1$ of size $n=10,000$ from noisy observations. (a) As a function of the noise variance for fixed $p=k=10$. (b) As a function of the number of samples for a fixed noise variance given by $\sigma^2_{\text{coeff}}=10^{-4}$

the adjacency matrix of the obtained graph. Let $\mathbf{x}\in\mathbb{R}^n$ be a k-bandlimited signal defined on the nodes of $\mathcal{G}$, then the frequency components are given by $\tilde{\mathbf{x}}_k$ such that $\mathbf{x}=\mathbf{V}_k\tilde{\mathbf{x}}_k$. In what follows, we model this setting as an inverse problem in which we have access to the graph signal $\mathbf{x}$ (the output) and we want to compute the GFT coefficients $\tilde{\mathbf{x}}_k$ (the input). The linear transform under study is $\mathbf{H}=\mathbf{V}_k$. The objective is to design a sketching matrix $\mathbf{H}_s\in\mathbb{R}^{k\times p}$ and a selection matrix $\mathbf{C}\in\mathcal{C}_{np}$ following the method described in Section IV, such that when applied to a stream of noisy measurements of the graph signal we obtain estimates of the GFT coefficients.

Denote by $\{\mathbf{x}_t+\mathbf{w}_t\}_{t=1}^{100}$ the stream of 100 realizations of the noisy output process. The graph process $\mathbf{x}$ is generated as a truly bandlimited signal $\mathbf{x}=\mathbf{V}_k\tilde{\mathbf{x}}_k$ where $\tilde{\mathbf{x}}_k$ is drawn at random as the absolute value of a standard normal random vector, i.e. $\tilde{\mathbf{x}}_k\sim|\mathbf{Z}|$ where $\mathbf{Z}\sim\mathcal{N}(0,\mathbf{I}_k)$. To model a more realistic setup, the covariance matrix $\mathbf{R}_x$ is not set to its ensemble value, but is estimated from a training set $\mathcal{T}$ consisting of $5,000$ noiseless realizations of $\mathbf{x}=\mathbf{V}_k\tilde{\mathbf{x}}_k$. The noise $\mathbf{w}$ has a normal distribution $\mathbf{w}\sim\mathcal{N}(0,\mathbf{R}_w)$ with $\mathbf{R}_w=\sigma^2\mathbf{I}_n$, independent of $\mathbf{x}$. The noise power σ^2 is considered known and proportional to the minimum energy of the signals in the training set $\mathcal{T}$, $\sigma^2=\sigma^2_{\text{coeff}}\cdot\min_{\mathbf{x}\in\mathcal{T}}\{\|\mathbf{x}\|_2^2\}$.

We consider a network of size $n=100$ where the signal is bandlimited with $k=10$ frequency components. The relative estimated MSE $\|\hat{\mathbf{y}}-\mathbf{y}\|_2^2/\|\mathbf{y}\|_2^2=\|\hat{\tilde{\mathbf{x}}}_k-\tilde{\mathbf{x}}_k\|_2^2/\|\tilde{\mathbf{x}}_k\|_2^2$ shown in Fig. 3 is obtained by averaging the errors computed in 500 realizations of the experiment. The estimator is $\hat{\tilde{\mathbf{x}}}_k=\mathbf{H}_s\mathbf{C}(\mathbf{x}+\mathbf{w})$ and $\mathbf{H}_s$ is given by (13). Several different sampling matrices $\mathbf{C}$ are considered. The curves in blue represent the results of using the sampling matrices obtained by the EDS method for different norms. The two curves in black correspond to the results of applying the convex relaxation heuristic described in Section V-A to solving problem (30). For the curve with the solid line and circle markings, the sampling vector is recovered by using tresholding whereas for the dotted line with square markings the random approach was adopted. The curve in red shows the results of applying the noise-blind heuristic explained in Section V-B to approximate the solution of problem (28). Finally, the curve in orange shows the outcome when the sampling matrix $\mathbf{C}$ is obtained by adopting the greedy approach (Section V-C). In addition, we consider solving the suboptimal problem of sampling rows (Section IV-B) by applying the convex relaxation heuristic and the greedy approach. The results are shown in the green curves.

In Fig. 3a, the estimated MSE is shown as a function of the noise coefficient σ^2_{coeff} for a fixed sampling size of $p=k=10$. First, it is clearly observed that, as the noise power increases, the performance of the estimators worsens since the measured value of each sample is less accurate. Then, we also observe that all three heuristics discussed in Section V to solve the inverse problem work better than the EDS approaches across all the range of noise power under consideration, demonstrating the merits of the design approach advocated in the paper. The relative improvement can be as high as one order of magnitude. The noise-blind heuristic arises as an alternative solution with low off-line computational cost achieving MSE within the range of $9\cdot10^{-5}$ and 10^{-3} for all noise power values simulated. It is also observed that it has better performance than the EDS strategies. Finally, note that applying convex relaxation heuristic and recovering the sampling vector by tresholding yields better results than using the random sampling recovery technique. With respect to the subptimal approach of sampling rows directly, it is observed that the resulting performance is not as good as any of the other methods. However, this strategy still yields an MSE of 10^{-2}. All in all, this simulation shows that, by using the greedy approach, we are able to obtain an estimate of the frequency components with a relative MSE of $2\cdot10^{-5}$ (for $\sigma^2_{\text{coeff}}=10^{-5}$) reducing the number of online operations by a factor of 10.

In Fig. 3b the simulation is carried out for varying number of samples selected p, for a fixed noise variance given by $\sigma^2_{\text{coeff}}=10^{-4}$. It is observed that when p grows and more samples are taken the estimation improves but this comes at the expense of a higher online computational cost incured by processing more samples. It is interesting to note that, for

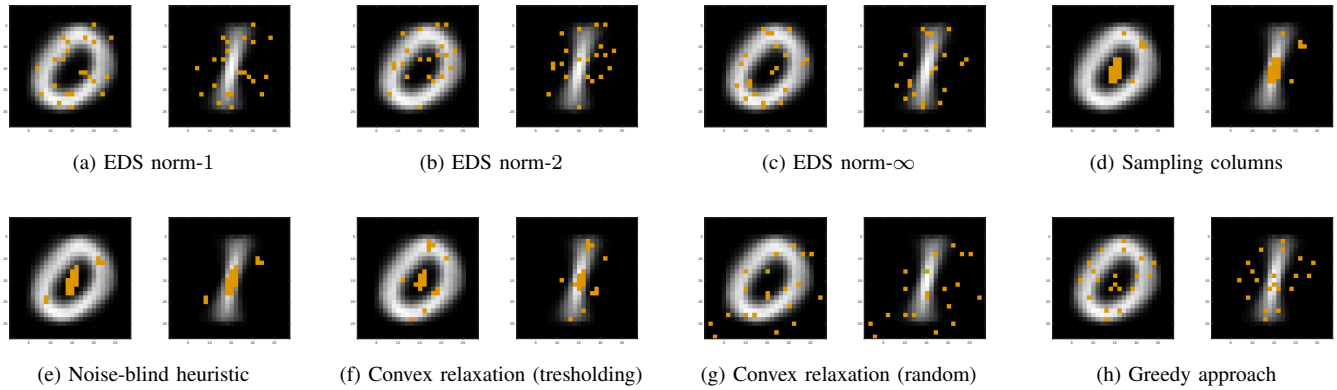

(a) EDS norm-1 (b) EDS norm-2 (c) EDS norm-∞ (d) Sampling columns

(e) Noise-blind heuristic (f) Convex relaxation (tresholding) (g) Convex relaxation (random) (h) Greedy approach

Figure 5. Selected pixels to use for classification of the digits according to each strategy. It is observed how the methods for reconstruction (a)–(c) tend to select pixels around the annulus that determine the digit 0, especially on top and bottom, which also help in reconstructing the digit 1. On the other hand, methods for classification (e)–(h) tend to distribute the pixel selection both in the center and in the annulus on the sides, which are the pixels that best help distinguish a zero (no pixels in the center) and a one (no pixels on the sides of the annulus). Last, the method of sampling columns (d) select most of the pixels in the center which is a reasonable choice, and some few pixels on the side.

$p < k$, the EDS sampling methods yield a better performance than using sketching-as-sampling strategies and for $p \geq k$ the latter methods show better results. As p grows, the noise-blind heuristic provides increasingly accurate approximations with low off-line computational cost. All things considered, it is observed that, by using the greedy approach with $p = 10$ samples, an estimate with an accuracy of 10^{-4} can be obtained. The online cost of obtaining such an estimate is 10 times lower than the one obtained by using all the available samples of each measurement. Also, when using $p = 24$ samples (and therefore reducing the online computational cost by a factor of 4.167) an estimate with an accuracy of $3{\cdot}10^{-5}$ can be obtained by employing either the greedy approach or the noise-blind heuristic.

B. Approximating the GFT as a direct problem

In this section we also consider the problem of estimating the frequency coefficients of a graph signal using only few samples, but we model this setting as a direct problem so as to illustrate both approaches. More specifically, we consider a k-bandlimited graph signal $\mathbf{x}$ with GFT coefficients given by $\tilde{\mathbf{x}}_k = \mathbf{V}_k^H \mathbf{x}$ so that $\mathbf{x}$ is regarded as the input, $\tilde{\mathbf{x}}_k$ as the output and $\mathbf{H} = \mathbf{V}_k^H$ as the linear transform to apply. Again, denote by $\{\mathbf{x}_t + \mathbf{w}_t\}_{t=1}^{100}$ a stream of 100 noisy realizations of the input process. The statistics of the process and the noise are the same as the ones described in Section VI-A.

For this section, we consider a substantially larger problem with an ER graph of $n = 10,000$ nodes and where edges connecting pairs of nodes are drawn independently with $p = 0.1$. The signal under study is bandlimited with $k = 10$ frequency coefficients. For these simulations we analyze results obtained when the sampling matrix $\mathbf{C}$ used is calculated by means of all three EDS sampling schemes (in blue), and also the results obtained when the sampling matrix is constructed by using the noise-blind heuristic (in red) and the greedy approach (in orange) to solve the problem in Section III. In addition, we consider the use of a sampling matrix $\mathbf{C}$ obtained by the suboptimal strategy of sampling columns (Section III-B), solved by employing a greedy approach (in green).

The relative MSE (estimated from 500 realizations of the experiment) obtained by using each method for a fixed number of samples $p = 10$ and varying noise power σ^2_{coeff} is depicted in Fig. 4a. Likewise, the relative MSE for the case of fixed noise variance given by $\sigma^2_{\text{coeff}} = 10^{-4}$ and varying number of samples p, is shown in Fig. 4b. We observe that, in both simulations, the heuristic approaches that approximate the solution to the binary optimization of the direct problem outperform the EDS sampling schemes. We also observe that, among the tested alternatives, the greedy approach always yields the best results. Indeed, when $\sigma^2_{\text{coeff}} = 10^{-4}$ and $p = 16$, the obtained MSE for the greedy strategy is $2 \cdot 10^{-2}$, while the online computational cost has been reduced by a factor of 625. Also, for $\sigma^2_{\text{coeff}} = 10^{-5}$ and $p = 10$ (a reduction by a factor of $1,000$ in the online computational cost) yields a relative MSE of $5 \cdot 10^{-3}$. It is also worth noticing that the relative MSE resulting from the sketching-as-sampling technique of sampling columns is at least one order of magnitude worse than the one obtained by the other sampling strategies, illustrating the relevance of the joint design of the sampling and the sketch.

C. MNIST handwritten digits classification

In this last example we consider the problem of classifying images of handwritten digits. This classification is usually carried out by applying a linear classifier to the coefficients obtained from a PCA transform of the image [37]. The transformation of images to the PCA domain is computationally expensive, and it is a major hindrance if each image in the stream has to be transformed prior to classification. Yet, observe that the PCA transform is a linear one and thus, it can be subsumed with the linear classifier into a single linear transform $\mathbf{H}$. Then, we can jointly design a smaller linear system $\mathbf{H}_s$ that operates only over a selected subset of pixels $\mathbf{Cx}$. This will markedly improve the online classification speed.

The images of handwritten digits are obtained from the MNIST database [35]. This database consists of images of

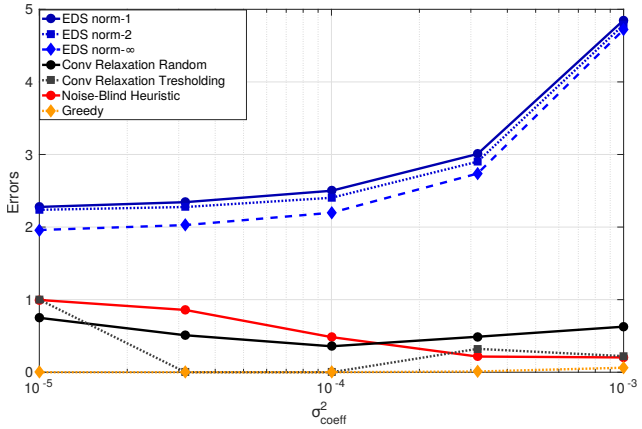

(a) Function of the noise variance

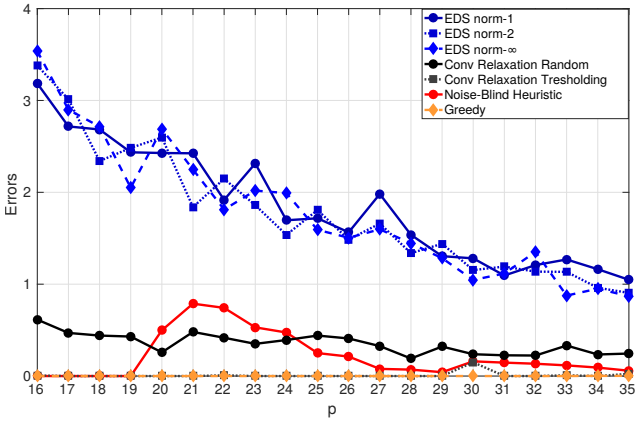

(b) Function of the sample size

Figure 6. Number of classification errors (out of 200 images to classify). (a) As a function of the noise variance coefficient σ^2_{coeff} for a fixed number of samples $p = k = 20$. (b) As a function of the number of samples p for a fixed noise variance given by $\sigma^2_{\text{coeff}} = 10^{-4}$.

size 28×28 portraying all ten handwritten digits. It contains a training set of $60,000$ images and a test set of $10,000$ images. We consider digits `0` and `1` and form a subset of the training set containing $5,000$ images of each digit drawn at random. Let $\mathcal{T} = \{\mathbf{x}_i\}_{i=1}^{10,000}$ denote such a set. The elements $\mathbf{x}_i$ are considered to be vectorized versions of the images. That is, the 28×28 images have been vectorized columnwise into a single vector $\mathbf{x} \in \mathbb{R}^n$ with $n = 784$. This training set $\mathcal{T}$ is used to estimate the covariance matrix $\mathbf{R}_x = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^H$ and this covariance matrix is used as the graph shift operator so that implementing PCA coincides with the notion of performing a GFT. For the linear classifier, we use a Support Vector Machine (SVM) [38]. The linear transform $\mathbf{H}$ to be applied is the cascade of the PCA transform $\mathbf{V}_k^H$ with the k largest components and the SVM classifier on this PCA domain. The overall dimension of $\mathbf{H}$ is of $m \times n$, with $m = 1$. The resulting scalar is used to determine which digit the image corresponds to. In all simulations, the number of PCA components retained is $k = 20$. For the streaming set $\{\mathbf{x}_t + \mathbf{w}_t\}_{t=1}^{200}$ we consider 100 vectorized images of each digit obtained at random from the test set provided in the MNIST database. We also add noise $\mathbf{w} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$ with known $\sigma^2 = \sigma^2_{\text{coeff}} \min_{\mathbf{x}_i \in \mathcal{T}} \{\|\mathbf{x}_i\|_2^2\}$, independent of $\mathbf{x}$.

For the first simulation, we set a fixed noise power given by $\sigma^2_{\text{coeff}} = 10^{-4}$ and $p = k = 20$ pixels sampled. In Fig. 5 the selected pixels for each strategy are illustrated. The error rate, measured as the number of misclassified images relative to the total of 200 images, for the SVM classifier using the full image is 0% (perfect classification). The error rates resulting from using a sampling matrix obtained by adopting the EDS sampling schemes are 0%, 4.81% and 1.2% for the EDS norm-1, norm-2 and norm-∞ respectively. When the sampling matrix is obtained by considering the different heuristic solutions to the direct sketching problem, the error rates resulting are 0% and 0.3% when adopting the convex relaxation technique (tresholding and random sampling recovery, respectively), 0.5% when employing the noise-blind heuristic and 0% when the greedy approach is used. Additionally, when using a sampling columns scheme, the resulting error rate is 2.71%. From the error rate analysis, we observe that, by using the sketching-as-sampling method proposed, we obtain the same classification performance as if the full image was used ($n = 784$), but using only $p = 20$ pixels, thereby reducing the online computational cost by a factor of 39.2.

In Fig. 5, we illustrate the sampled pixels according to the selection matrix obtained by the different methods under study. From the averaged images of digits `0` and `1` we note that the most useful pixels for classification are those that are white in one image and black in the other, and viceversa. For instance, the pixels located in the center of the image are black for a `0` and white for a `1`. Likewise, pixels on the left and right of the center (i.e. those pixels situated on the annulus of the `0` both right and left of the center) are white for digit `0` and black for a `1`. Intuitively, these pixels should be the most useful ones in discriminating between digits, and therefore, should be selected by the proposed sketching-as-sampling method. We observe in Figs. 5e–h that this holds true, especially for pixels selected by employing the greedy approach and the noise-blind heuristic, where it is evident that these selected pixels are mostly white in one image and black in the other, and viceversa. On the other hand, the EDS techniques that aim at reconstructing the image, tend to select pixels in the center of the image but also above and below it, as observed in Figs. 5a–c. We note that pixels above and below the center are shared by both digits and hence are useful in reconstructing the image (since they are present in both images). Finally, we observe in Fig. 5d that that the sampling columns scheme selects the bulk of the pixels in the center of the image. The relatively high error rate might be explained by the inability of the matrix sketch (constructed by a subset of columns of the full linear transform) to get rid of the noise.

For the last two simulations we analyze the number of misclassified images as a function of the noise power for a fixed number of samples $p = k = 20$, see Fig. 6a; and as a function of the number of samples for a fixed noise variance given by $\sigma^2_{\text{coeff}} = 10^{-4}$, see Fig. 6b. In both cases, we observe that the selection matrix built by the heuristics proposed in Section V outperform the ones built by employing EDS techniques. Moreover, the proposed sketching-as-sampling methods average less than one classification error.

VII. Conclusions

A problem setup where a set of streaming signals with large dimensionality must be processed using a linear operator was investigated. Assuming that the input signals lie in a low-dimensional subspace, the goal was to develop optimal sampling and sketching schemes to reduce the online processing burden. The sampling scheme selects the subset of input values to be processed, while the sketching scheme specifies the reduced-dimensionality transformation to be applied to the sampled signals. Different from traditional (sampling) schemes aimed at reducing the distortion between the sampled and the original *input* signal, the objective here was to obtain an *output* as close as possible to the one generated by the original (large) linear operator. The joint design problem was formulated as a two-stage optimization where we first found the expression for the optimal sketching matrix as a function of the sampling strategy and then substituted this expression into the original cost to solve for the optimal sampling scheme. Since the resultant sampling problem was a non-convex binary optimization, different suboptimal schemes were proposed and their complexity was discussed. Numerical tests using synthetic data as well as real data for handwritten digit classification demonstrated the merits of the proposed approach.

References

[1] F. Gama, A. G. Marques, G. Mateos, and A. Ribeiro, "Rethinking sketching as sampling: Linear transforms of graph signals," in *50th Asilomar Conf. Signals, Systems and Computers*, Pacific Grove, CA, 6-9 Nov. 2016.

[2] ——, "Rethinking sketching as sampling: Efficient approximate solution to linear inverse problems," in *2016 IEEE Global Conf. Signal and Inform. Process.* Washington, DC: IEEE, 7-9 Dec. 2016.

[3] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 83–98, May 2013.

[4] A. Sandryhaila and J. M. F. Moura, "Big data analysis with signal processing on graphs: Representation and processing of massive data sets with irregular structure," *IEEE Signal Process. Mag.*, vol. 31, no. 5, pp. 80–90, Sep. 2014.

[5] S. K. Narang and A. Ortega, "Perfect reconstruction two-channel wavelet filter banks for graph structured data," *IEEE Trans. Signal Process.*, vol. 60, no. 6, pp. 2786–2799, June 2012.

[6] A. Sandyhaila and J. M. F. Moura, "Discrete signal processing on graphs: Frequency analysis," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3042–3054, June 2014.

[7] K. Slavakis, G. B. Giannakis, and G. Mateos, "Modeling and optimization for big data analytics: (statistical) learning tools for our era of data deluge," *IEEE Signal Process. Mag.*, vol. 31, no. 5, pp. 18–31, Sep. 2014.

[8] D. K. Berberidis, V. Kekatos, and G. B. Giannakis, "Online censoring for large-scale regressions with application to streaming big data," *IEEE Trans. Signal Process.*, vol. 64, no. 15, pp. 3854–3867, Aug. 2016.

[9] M. W. Mahoney, "Randomized algorithms for matrices and data," *Foundations and Trends® in Machine Learning*, vol. 3, no. 2, pp. 123–224, 2011.

[10] D. P. Woodruff, "Sketching as a tool for numerical linear algebra," *Foundations and Trends® in Theoretical Computer Science*, vol. 10, no. 1-2, pp. 1–157, 2014.

[11] S. Chen, R. Varma, A. Sandryhaila, and J. Kovačević, "Discrete signal processing on graphs: Sampling theory," *IEEE Trans. Signal Process.*, vol. 63, no. 24, pp. 6510–6523, Dec. 2015.

[12] A. G. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Sampling of graph signals with successive local aggregations," *IEEE Trans. Signal Process.*, vol. 64, no. 7, pp. 1832–1843, Apr. 2016.

[13] M. Tsitsvero, S. Barbarossa, and P. Di Lorenzo, "Signals on graphs: Uncertainty principle and sampling," *arXiv:1507.08822v3 [cs.DM]*, 17 May 2016. [Online]. Available: http://arxiv.org/abs/1507.08822

[14] R. M. Gray, "Quantization in task-driven sensing and distributed processing," in *31st IEEE Int. Conf. Acoust., Speech and Signal Process.* Toulouse, France: IEEE, 14-19 May 2006.

[15] S. Liu, S. P. Chepuri, M. Fardad, E. Masazade, G. Leus, and P. K. Varshney, "Sensor selection for estimation with correlated measurement noise," *IEEE Trans. Signal Process.*, vol. 64, no. 13, pp. 3509–3522, July 2016.

[16] S. P. Chepuri and G. Leus, "Subsampling for graph power spectrum estimation," in *2016 IEEE Sensor Array Multichannel Signal Process. Workshop*. Rio de Janeiro, Brazil: IEEE, 10-13 July 2016.

[17] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.

[18] M. Püschel and J. M. F. Moura, "Algebraic signal processing theory: Foundation and 1-d time," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3575–3585, Aug. 2008.

[19] F. R. K. Chung, *Spectral Graph Theory*, ser. Conference Board of the Mathematical Sciences: Regional Conference Series in Mathematics. Providence, RI: American Mathematical Society, 1997, no. 92.

[20] B. Girault, P. Gonçalves, and E. Fleury, "Translation and Stationarity for Graph Signals," École Normale Supérieure de Lyon, Inria Rhône-Alpes, Research Report RR-8719, Apr. 2015.

[21] A. G. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Stationary graph processes and spectral estimation," *arXiv:1603.04667v1 [cs.SY]*, 14 March 2016. [Online]. Available: http://arxiv.org/abs/1603.04667

[22] N. Perraudin and P. Vandergheynst, "Stationary signal processing on graphs," *arXiv:1601.02522v4 [cs.DS]*, 8 July 2016. [Online]. Available: https://arxiv.org/abs/1601.02522

[23] S. Segarra, A. G. Marques, G. Leus, and A. Ribeiro, "Space-shift sampling of graph signals," in *41st IEEE Int. Conf. Acoust., Speech and Signal Process.* Shanghai, China: IEEE, 21-25 March 2016.

[24] S. Segarra, A. G. Marques, G. Mateos, and A. Ribeiro, "Network topology inference from spectral templates," *arXiv:1608.03008v1 [cs.SI]*, 10 Aug. 2016. [Online]. Available: http://arxiv.org/abs/1608.03008

[25] B. Pasdeloup, V. Gripon, G. Mercier, D. Pastor, and M. G. Rabbat, "Characterization and inference of weighted graph topologies from observations of diffused signals," *arXiv:1605.025691v1 [cs.DS]*, 9 May 2016. [Online]. Available: http://arxiv.org/abs/1605.02569

[26] M. Mardani, G. Mateos, and G. B. Giannakis, "Subspace learning and imputation for streaming big data matrices and tensors," *IEEE Trans. Signal Process.*, vol. 63, no. 10, pp. 2663–2677, Nov. 2015.

[27] G. H. Golub and C. F. van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD: Johns Hopkins University Press, 1996.

[28] W. B. Johnson and J. Lindenstrauss, "Extensions of lipschitz mappings into a hilbert space," *Contemporary Mathematics*, vol. 26, p. 189206, Jan. 1984.

[29] S. Segarra, A. G. Marques, G. Leus, and A. Ribeiro, "Reconstruction of graph signals through percolation from seeding nodes," *IEEE Trans. Signal Process.*, vol. 64, no. 16, pp. 4363–4378, Aug. 2016.

[30] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, UK: Cambridge University Press, 1985.

[31] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.

[32] G. Puy, N. Tremblay, R. Gribonval, and P. Vandergheynst, "Random sampling of bandlimited graph signals," *arXiv:1511.05118v2 [cs.SI]*, 20 May 2016. [Online]. Available: http://arxiv.org/abs/1511.05118

[33] D. E. Knuth, *The Art of Computer Programming*, 2nd ed. Upper Saddle River, NJ: Addison-Wesley, 1998, vol. 3: Sorting and Searching.

[34] L. F. O. Chamon and A. Ribeiro, "Near-optimality of greedy set selection in the sampling of graph signals," in *2016 IEEE Global Conf. Signal and Inform. Process.* Washington, DC: IEEE, 7-9 Dec. 2016.

[35] Y. Le Cun, C. Cortes, and C. J. C. Burges, "The MNIST database of handwritten digits," Website, 18 Aug. 2015. [Online]. Available: http://yann.lecun.com/exdb/mnist/

[36] R. Varma, S. Chen, and J. Kovačević, "Spectrum-blind signal recovery on graphs," in *2015 IEEE Int. Workshop Comput. Advances Multi-Sensor Adaptive Process.* Cancún, México: IEEE, 13-16 Dec. 2015, pp. 81–84.

[37] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. Upper Saddle River, NJ: Pearson Prentice Hall, 2008.

[38] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York, NY: John Wiley and Sons, 2001.